并且在道路的左手侧驾驶的国家中位于前座的右侧。美国国家公路交通安全管理局推荐为了安全原因驾驶员自身与驾驶员安全气囊相距10英尺(25.4厘米)。该距离从方向盘的中心到驾驶员的胸骨测得。因为车辆200是不透明的,车辆200的部分(例如,车顶、车门、挡风玻璃框等等)可以使始于典型驾驶员定位的视角模糊。另外,虽然人类具有几乎180度的面向前方的水平视场,人类视场的垂直范围明显较小。然而,驾驶员辅助系统可以使用具有超过坐在典型驾驶员定位上的驾驶员的视角的延伸视角的图像捕获装置(例如,图像捕获装置122)。延伸视角可以使得系统100识别驾驶员无法检测的在车辆200的环境中的对象。而且,系统100可以确定与所识别的对象相关联的变化,并因此进行动作。例如,系统100可以确定位于车辆200上方的区域中的交通灯的状态变化。作为另一示例,系统100可以确定位于车辆200旁边的区域中的不同对象(例如,行人)的轨迹。

[0249] 图16和图17是与公开的实施方式相一致的在两种情况下的车辆200的概略表示。如图16和图17中所例示的,车辆200可以包括与第一视角相关联的图像捕获装置122。第一视角在垂直方向上由FOV 1610来表示,如图16中所示,并且在水平方向上由FOV 1710来表示,如图17中所示。图16和图17中同样例示的是始于车辆200中的典型驾驶员定位的第二视角。第二视角在垂直方向上由FOV 1620来表示,如图16中所示,并且在水平方向上由FOV 1720来表示,如图17中所示。与本公开的实施方式相一致,与图像捕获装置122相关联的第一视角可以大于与车辆200中的典型驾驶员定位相关联的第二视角。较大的第一视角与第二视角之间的增量在本公开中被标记为延伸区域。在一些实施方式中,延伸区域可以包括系统100可以识别的位于车辆200中的典型驾驶员定位的视线的外部的对象。另外,系统100可以基于与在延伸区域中识别的对象相关联的所确定的变化来触发一个或更多个动作(例如,产生导航响应)。

[0250] 图16例示了用于确定与位于车辆200的典型驾驶员定位的视线的外部的一个或更多个对象相关联的变化的第一示例。图16中使用FOV 1610和FOV 1620的角范围,用于例示目的。在一些实施方式中,FOV 1620可以是例如大约65度,FOV 1610可以是例如高于65度。在其它实施方式中,FOV 1620可以是例如大约135度,FOV 1610可以是例如高于150度、高于160度或者高于170度。在一些实施方式中,FOV 1610和FOV 1620可以在从水平面开始的向上方向上测得。如图16中所示,延伸区域可以包括位于车辆200上方的区域中的交通灯架1630。位于典型驾驶员定位的视线外部的交通灯架1630显示红灯。在一些实施方式中,系统100能够在车辆200的驾驶员甚至无法看到交通灯架的情况下检测交通灯和交通灯的显示颜色的变化。

[0251] 这些情况的一个示例是车辆200的驾驶员以促使车辆200停在十字路口的停止线和/或将车辆200放置为太靠近交通灯架1360的方式施加制动。这些情况的另一示例是车辆200到达十字路口处的停止线,并且交通灯架1630被放置在高杆上所以其能够从远处被看到。因为交通灯架1630位于FOV 1620的外部,当交通灯架1630中的灯变绿时,坐在典型驾驶员定位上的驾驶员将不能检测状态变化。然而,因为交通灯架1630位于FOV内部并因此考虑到图像捕获装置122,系统100将能够检测交通灯的状态变化。识别位于车辆200上方的区域中的交通灯可以提高安全性和驾驶员便利性。例如,当考虑车辆200的驾驶员和乘客的安全性时,如果交通灯架1630不在FOV 1620内,则使车辆200倒退直至交通灯架1630位于FOV 1620内可能是不安全的。另外,即使使车辆200倒退不危险,对驾驶员来说使车辆200倒退也

IPR2025-00943
Tesla EX1002 Page 968

可能是不方便的。在一个实施方式中，在确定交通灯架1630中的灯已经变绿之后，系统100可以促使车辆200加速以通过十字路口。

[0252] 图17例示了用于确定与位于车辆200的典型驾驶员定位的视线的外部的对象相关联的变化的第二示例。图17中使用FOV 1710和FOV 1720的角范围，用于例示目的。在一些实施方式中，FOV 1720可以是例如大约180度，FOV 1710可以是例如高于200度、高于240度或者高于300度。如图17中所示，延伸区域可以包括位于车辆200旁边的区域中的行人1730。行人1730可以在特定方向上移动。在一些情况下，行人1730可以被定位在FOV 1720的边缘，但是因为驾驶员集中于前方的道路上，其可能无法检测行人1730。相反，系统100不但具有较大的FOV，其不必集中于前方的道路上并且能够识别位于其FOV的任何区域中的对象。因为行人213位于FOV 1710内，系统100可以检测行人1730并且确定行人的行走方向轨迹。在其它实施方式中，系统100可以识别位于车辆200的旁边的对象、确定对象的移动方向轨迹、确定对象是否在具有车辆200的可能碰撞路线上、并因此进行动作。例如，在确定行人1730在具有车辆200的可能碰撞路线上之后，系统100可以控制车辆的喇叭以生成一种类型的音频警报。

[0253] 如上所讨论的，图3B至图3D例示了被配置为被定位在后视镜310的后面并且靠着车辆挡风玻璃的示例性相机安装件370。如图3B至图3D中所示，相机安装件370可包括多个图像捕获装置(122、124和126)。在一个实施方式中，图像捕获装置122、124和126的组合的视角形成第二视角。例如，为了使FOV 1710高于300度，可以使用六个图像捕获装置的组合。在不同实施方式中，图3B至图3D中描述的图像捕获装置中的至少一个图像捕获装置(例如，图像捕获装置122)包括超宽透镜。超宽透镜可以使得图像捕获装置122能够具有第一视角。在一些实施方式中，超宽透镜可以具有比图像捕获装置122的图像传感器的短边短的焦距。另选地，超宽透镜可以具有比图像传感器的长边短、但是比图像传感器的短边长的焦距。与本公开的实施方式相一致，图像捕获装置122可以被配置为捕获相对于图像捕获装置122的定位被定位在例如150度到180度之间、150度到190度之间、160度到180度之间或者170度到180度之间的圆锥曲线区域中的交通灯的图像。

[0254] 图18是与公开的实施方式相一致的存储器的示例性框图，该存储器被构造成存储用于执行一个或多个操作的指令。如图18中所示，存储器140或150可存储对象识别模块1802、状态确定模块1804和系统响应模块1806。公开的实施方式并不限于任何特定构造的存储器140或150。进一步，处理单元110可执行存储在任何模块1802-1806(包括在存储器140或150中)中的指令。

[0255] 在一个实施方式中，对象识别模块1802可存储软件指令(诸如计算机视觉软件)，当由处理单元110执行时，该软件指令识别位于车辆200中的典型驾驶员定位的视线外部的对象。在一个实施方式中，处理单元110可以扫描多个图像，将图像的部分与一个或更多个预定图案进行比较，并且识别在车辆200旁边的区域中的交通灯。在另一实施方式中，处理单元110可以扫描多个图像，将图像的部分与一个或更多个预定图案进行比较，并且识别位于车辆200旁边的区域中的对象。在一些实施方式中，模块1802可以存储软件指令，当由处理单元110执行时，该软件指令识别十字路口并确定交通灯可能在附近。处理单元110识别十字路口的一种方式是通过使用历史图像数据。例如，处理单元110可以在即使当在当前图像中看不到交通灯、但是在之前图像(例如，前一帧、五帧之前等等)中曾看到交通灯时也可

IPR2025-00943
Tesla EX1002 Page 969

以确定交通灯可能在附近。另外地或另选地，处理单元110可以识别指示附近十字路口的图像数据对象。例如，处理单元110可以通过识别前方道路的图像数据中的十字交义路来确定交通灯可能在附近。在另一示例中，处理单元110可以在图像数据中识别控制在与车辆200的行驶方向相反的方向上进行的交通的其它交通灯的后部。另外地或另选地，处理单元110可以使用地图数据(诸如来自地图数据库160的数据)来识别十字路口。例如，处理单元110可以考虑从地图数据库160中提取的通常道路形状描述，以确定交通灯可能在附近。

[0256] 在另一实施方式中，模块1802可以存储软件指令，当由处理单元110执行时，该软件指令对多个图像中的光学畸变进行校正以使得能够识别对象。光学畸变可以是图像捕获装置122使用超宽透镜的结果。例如，光学畸变可以是图像放大率随着距光学轴的距离而减小的桶形畸变，并且因此该畸变可以最接近图像帧的边缘。与本公开相一致，可以使用多种算法来校正可以使用超宽透镜产生的光学畸变。例如，可以使用一种算法来校正光学畸变，以识别接近图像帧的边缘的对象。作为另一示例，可以使用不同算法来识别位于图像帧的边缘处的对象的状态。在另一实施方式中，图像捕获装置可以包括多个图像捕获装置(例如，图像捕获装置122、124和126)，并且每个图像捕获装置可以与不同视场相关联。模块1802可以存储软件指令，当由处理单元110执行时，该软件指令可以对从多个图像捕获装置中导出的图像正确进行组合。例如，即使当在两个图像捕获装置的视场中存在一些重叠时，对象的不同部分也可以由不同的图像捕获装置来捕获。因此，对象识别模块1802可以包括用于识别跨越由不同图像捕获装置捕获的相邻图像的区域中的对象的指令。

[0257] 在一个实施方式中，状态确定模块1804可以存储软件指令，当由处理单元110执行时，该软件指令确定在图像数据中识别并且至少部分地位于延伸区域内部的对象的状态。与本公开相一致，延伸区域包括图像捕获装置的视场(例如，FOV 1610和1710)与车辆200的典型驾驶员定位的视场(例如，FOV 1620和1720)之间的区域。在一个实施方式中，状态确定模块1804可存储软件指令，当由处理单元110执行时，该软件指令确定位于车辆200上方的区域中的交通灯的状态。例如，处理单元110可以通过将一个或更多个图像与所存储的图案进行比较来确定交通灯的状态。另外，可以基于所存储的图案与在不同的时间拍摄的两个或更多个图像之间的差来确定交通灯的状态的变化。在另一实施方式中，状态确定模块1804可存储软件指令，当由处理单元110执行时，该软件指令确定位于车辆200旁边的区域中的移动对象的轨迹。例如，处理单元110可以通过将多个图像彼此进行比较来确定移动对象的轨迹。另外，处理单元110可以确定移动对象是否位于具有车辆200的可能碰撞路线上。

[0258] 在一个实施方式中，系统响应模块1806可存储软件指令，当由处理单元110执行时，该软件指令在车辆200中产生一个或更多个响应。一种类型的系统响应可以包括导航响应，如上面参照导航响应模块408详细描述的。其它类型的系统响应可以涉及控制节流系统220、制动系统230和/或转向系统240。例如，处理单元110可发送电子信号，致使系统100物理地踩下制动器达预定的量或者部分地松开车辆200的加速器。进一步，处理单元110可发送致使系统100车辆200沿特定方向转向的电子信号。这种响应可以基于关联交通灯的确定。其它系统响应可以包括提供各种通知(例如，警告和/或警报)。通知可以包括：报告位于车辆200上方的交通灯的状态；报告位于车辆200上方的交通灯的状态的变化；报告在车辆旁边存在对象；以及报告对象位于具有车辆200的碰撞路线上。

[0259] 图19是示出用于基于如图16中描述的位于车辆200上方的区域中的交通灯的状态

IPR2025-00943
Tesla EX1002 Page 970

的变化来对车辆进行导航的示例性处理1900。预期一种用于基于如图17中所示的位于车辆200的旁边的区域中的移动对象的轨迹来对车辆进行控制的相似处理。因此，示例性处理1900可以适用于各种情况，包括图17中描述的情况。

[0260] 如图19中所示，在步骤1902，图像捕获装置122可以在典型驾驶员定位的视线的外部获取位于车辆200上方的区域中的交通灯的多个图像。在步骤1904，处理单元110可以基于车辆位置的至少一个指示符(例如，车道标记)来确定多个图像中的所捕获的交通灯的关联。在一些实施方式中，该步骤可以是可选的并且可省略。在步骤1906，处理单元110可以基于来自多个图像当中的至少第一图像的分析来确定交通灯的状态。在步骤1908，处理单元110可以基于来自多个图像当中的至少第二图像来确定交通灯的状态的变化。例如，交通灯的状态的变化可以是从红灯变为绿灯。在确定交通灯的状态已经变化之后，在步骤1910，处理单元110可以基于所确定的状态变化来产生系统响应。处理1900的这些步骤下面将更加详细地进行讨论。

[0261] 在步骤1902，图像捕获装置122可以获取位于车辆200上方的区域中的交通灯的多个图像。与公开的实施方式相一致，图像捕获装置122可以包括一个或更多个图像捕获装置。与一些实施方式相一致，交通灯可以位于车辆200中的典型驾驶员定位的视线的外部。例如，交通灯可以位于如参照图16进行描述的延伸区域中。在一些实施方式中，图像捕获装置122可位于车辆200的内部。另外，虽然在一些实施方式中图像捕获装置122可以位于后视镜后面并且靠着车辆挡风玻璃，但是处理单元110可以通过挡风玻璃来确定交通灯的状态。

[0262] 在步骤1904，处理单元110可以基于车辆位置的至少一个指示符来确定位于车辆200上方的区域中的交通灯的关联。系统100如何可以确定交通灯的关联的详细公开将在下面结合图20至图23来给出。在一些实施方式中，车辆位置的指示符可以从对多个图像的分析中导出。例如，处理单元110可以基于多个图像中的车道标记的识别来确定交通灯与车辆200的相对位置。另外地或另选的，车辆位置的指示符可以由从位置传感器130中检索出的地理位置数据中导出。例如，处理单元110可以基于GPS信号或当地定位信号来确定交通灯与车辆200的相对位置。在一些实施方式中，处理单元110可以将GPS获取的车辆位置与地图数据进行比较，以确定交通灯的关联。例如，处理单元110可以使用从图像数据和地图数据中导出的信息来进行分析。使用所导出的信息，处理单元110可以确定位于车辆200上方的区域中的所检测的交通灯与车辆200当前正在行驶的车道之间的对应。

[0263] 在步骤1906，处理单元110可以基于来自多个图像当中的至少第一图像的分析来确定交通灯的状态。在一些实施方式中，位于车辆200上方的区域中的交通灯的状态与交通灯指示的颜色相关联。例如，稳定的绿灯可以指示车辆200在给道路上的其它车辆、自行车或行人让路之后能够继续；闪烁的黄灯是警告并且可以指示车辆200应当谨慎行事；并且稳定的红灯可以指示车辆200需要停止。另外地或另选地，交通灯的状态还可以包括与交通灯相关联的附加信息。该附加信息可以是交通灯的类型。例如，位于车辆200上方的区域中的交通灯可以是行人通过信号。因此，处理单元110可以确定行人可能通过车辆200前面的道路。在一些实施方式中，为了确定位于车辆200上方的区域中的交通灯的状态，处理装置110可以校正至少第一图像中的光学畸变。当图像捕获装置122包括超宽透镜时，光学畸变可以发生。

[0264] 在步骤1908，处理单元110可以基于来自多个图像当中的至少第二图像来确定交

IPR2025-00943
Tesla EX1002 Page 971

通灯的状态的变化。在一些实施方式中，位于车辆200上方的区域中的交通灯的状态可以例如从红灯变为绿灯。在一个实施方式中，至少第二图像可以包括继至少第一图像之后的两个或更多个图像。另外，处理单元110可以仅处理至少第二图像的一部分。例如，当车辆200不移动时，框架中交通灯的位置保持相同。因此，在一些实施方式中，于在至少第一图像中识别包括交通灯架的框架的一部分之后，处理单元110可以通过仅处理至少第二图像中的框架的相同部分来确定状态的变化。例如，在一些实施方式中，通过仅处理包括交通灯架的框架的一部分，处理单元110可以节省处理电力。

[0265] 在步骤1910，处理单元110可以基于所确定的状态变化来产生系统响应。在一些实施方式中，处理单元110可以产生一个或更多个系统响应(例如，两个或更多个响应)，包括不同类型的响应。一种类型的系统响应可以包括向车辆200的驾驶员提供各种通知(例如，警告和/或警报)。该通知可以经由扬声器360或者触摸屏320来提供。例如，系统响应可以包括将所述状态的变化报告给所述车辆200的驾驶员。另一种类型的系统响应可以包括导航响应，导航响应可以包括例如加速至驾驶员指定的巡航速度。

[0266] 识别关联交通灯

[0267] 在一些实施方式中，系统100可以在关联和非关联(或较少关联)交通灯之间进行区分。在典型的驾驶过程期间，车辆200可以通过一个或更多个枢纽(诸如具有多个交通灯的十字路口)。例如，十字路口处的一个或更多个交通灯可以控制在特定方向上朝向十字路口行驶的交通。因此，一个或更多个交通灯可以控制车辆200是否可以继续行驶通过十字路口或者车辆200是否必须停在十字路口处。然而，除了控制车辆200正在行驶的车道的交通灯以外，控制其他车道上的交通的其它交通灯对车辆200也是可视的。根据控制除了车辆200正在行驶的车道之外的车道的交通灯中的任何一个交通灯来对车辆200进行导航可能产生对车辆200的计划路线来说不合需要或不适用的导航响应。因此，为了使得能够以适于车辆200的计划导航路径的方式来进行车辆200的自助控制，系统100可以识别多个交通灯中的哪一个正在控制车辆200正在行驶的车道的交通同时忽略(或者较少关注)控制其它交通车道的其它交通灯。进一步，在系统100识别关联交通灯之后，系统100可以识别交通灯的状态(例如，红色、黄色、绿色)并实施合适的导航响应。例如，当对车辆200正在行驶的车道进行控制的红灯被识别出时，或者当对车辆200正在行驶的车道进行控制的黄灯被识别出且车辆200越过枢纽的预定距离时，系统100可以中止巡航控制并且施加制动。

[0268] 在道路上的关联与非关联(或较少关联)交通灯之间进行区分是复杂的。图20和图21例示了基于从不适用交通灯导出的信息的驾驶可能是不合需要的情况的示例。在图20中，车辆200a正在多车道道路上行驶。道路的每个车道与不同的交通灯架相关联。车辆200a正在接近十字路口并且在指定为前进通过十字路口并到达十字路口的另一侧的车道上行驶。同样如图20中所示，车辆200b在允许交通继续直行并通过十字路口或者允许交通进行右转弯的车道上行驶。与车辆200b正在行驶的车道相关联的交通灯架包括右转弯交通灯。作为另一示例，在图21中，车辆200a到达非垂直道路的枢纽。包括不对车辆200a正在行驶的车道进行控制的一些交通灯架的多个交通灯架(例如，交通灯架2122和2124)由于枢纽的定位而对车辆200a可见。

[0269] 返回图20，示出了具有以下驾驶选项的十字路口2000：具有通向道路B、D和F的车道的道路A；具有通向道路D和F的车道的道路C；具有通向道路F、H和B的车道的道路E；以及

IPR2025-00943
Tesla EX1002 Page 972

具有通向道路H、B和D的车道的道路G。道路A与四个交通灯架2012、2014、2016和2018相关联。在所示的情况下，每个交通灯架对不同车道进行控制。道路C与两个个交通灯架2022和2024相关联。交通灯架2024包括用于继续直行的交通灯2024a和用于右转弯的交通灯2024b(例如，当右转弯得到授权时，显示绿色箭头)。道路E与三个交通灯架2032、2034和2036相关联。并且道路G与两个个交通灯架2042和2044相关联。交通灯架2042包括用于继续直行的交通灯2042a和用于左转弯的交通灯2042b。在图20中例示的情况下，交通灯2042b和交通灯架2012、2018和2032显示绿灯，而全部的其它交通灯架显示红灯。当然，除了图20中示出的这些以外，还可以存在很多其它道路变型和相关交通灯配置。

[0270] 在图20中例示的情况下，车辆200a位于道路A上的继续直行通过十字路口2000到达道路D的车道上。然而，包括在车辆200a中的图像捕获装置的视场可以包括交通灯架2012和交通灯架2014双方(或者甚至其它架)。车辆200a在交通灯架2012和2014都显示红灯时到达十字路口2000，并且仅最近地交通灯架2012中的灯已经变绿。在这种情况下，重要的是车辆200a的交通灯检测系统100识别出架2012的绿灯不适用于车辆200a。但是，系统100应当使所确定的导航响应基于与另外的关联架2014相关联的灯的状态。

[0271] 在图20中例示的情况的另一方面，车辆200b位于道路C上的继续直行通过十字路口2000到达道路F的车道中，并且该车道还允许交通进行右转弯到达道路D。车辆200b面向包括用于继续直行的交通灯2024a和用于右转弯的交通灯2024b的交通灯架2024。车辆200b在交通灯架2024显示用于继续直行(交通灯2024a)且用于右转弯(交通灯2024b)的红灯时到达十字路口2000，并且仅最近地交通灯2024b已经变绿。这意味着交通灯架2024的当前状态阻止交通直行行驶和右转弯。如果车辆200b针对灯2024b的状态来动作而没有识别并考虑灯2024a的状态的话，则可能出现不期望的情况。例如，如果车辆200b基于交通灯2024b的信息(示出绿灯)而行驶到道路F(即，直行行驶)，则从道路E行驶到道路B的车辆可能产生对车辆200b的危险。

[0272] 转向图21，图21例示了具有具备以下行驶选项的非垂直道路的十字路口2100：具有两个车道的道路A，一个车道通向道路C，另一个车道通向道路E；具有通向道路C的两个车道以及还能够通向道路E的一个车道的道路B；以及具有通向道路E的单车道的道路D。道路A与两个交通灯架2112和211相关联，并且每个交通灯架对不同车道进行控制。道路B也与两个交通灯架2122和2124相关联，并且每个交通灯架对不同车道进行控制。道路D与单个交通灯架2132相关联。在图21中例示的情况下，交通灯架2122显示绿灯，而全部的其它交通灯架显示红灯。如图21中所示，车辆200a位于道路A上的继续直行通过十字路口2000到达道路C的车道上。然而，包括在车辆200a中的图像捕获装置122的视场可以包括交通灯架2114和交通灯架2122双方。在该示例中，车辆200a在交通灯架2114和2122都为红时到达十字路口2100，并且仅最近地交通灯架2122中的灯已经变绿。在这种情况下，重要的是，车辆200a基于交通灯架2114的状态而不是基于交通灯架2122的状态来确定其导航响应。

[0273] 图20和图21中描述的三种情况仅提供了道路状况的一些示例，其中，具有能够在关联与非关联交通灯架之间进行区分的驾驶员辅助系统可能是有帮助的，以确定包括在关联交通灯架中的交通灯的状态，并且基于该状态生成并作出合适的导航响应。进一步，这些示例证明系统可能需要对面向车辆的多个交通灯架进行评估，以便识别最适于车辆正在行驶的车道或最适于车辆的计划行驶方向的交通灯架(尤其是在例如多个交通灯或交通灯架

IPR2025-00943
Tesla EX1002 Page 973

可以与车辆正在行驶的单个车道相关联的情况下)。

[0274] 如上所讨论的,系统100可以在关联与非关联交通灯架之间进行区分,确定包括在(多个)关联交通灯架中的交通灯的状态,并且基于各种行驶场景中的状态来生成并作出合适的导航响应。例如,当车辆200接近十字路口时,系统100可以确定哪个交通灯是关联的,确定该交通灯的状态,并且发现由一个或更多个图像捕获装置122-126捕获的图像中的任何其它关联信息。如果交通灯为红,则系统100可以促使车辆200施加其制动。如果交通灯为绿,则系统100可以促使车辆200继续。如果交通灯为黄,则系统100可以基于图像的分析、车辆200的速度和/或位置数据(例如,GPS数据)来确定距十字路口的距离和/或估计到十字路口的时间。如果车辆200在预定时间(例如,五秒钟、十秒钟等等)和/或距离(例如,一米、五米、十米等等)阈值之内,则系统100可以促使车辆200继续。如果车辆200不在预定时间阈值和/或距离阈值之内,则系统100可以促使车辆200在到达十字路口之前停止。作为另一示例,当车辆200在交通灯处停止时,在交通灯将其状态从红变为绿之后,系统100可以产生导航响应,该导航响应例如包括施加加速器、释放制动器并且转向通过十字路口。

[0275] 图22是与公开的实施方式相一致的存储器140/150的示例性框图,该存储器140/150被构造成存储用于执行一个或多个操作的指令。存储器140/150可以被处理单元110访问。如上所讨论的,处理单元110可包括各种装置,诸如控制器、图像预处理器、中央处理单元(CPU)、支持电路、数字信号处理器、集成电路、存储器或任何其它类型的用于图像处理和分析的装置。因此,存储器140/150可以被处理单元110访问,可以与处理单元110集成,或者可以与处理单元110一起包括在嵌入式系统中。如图22中所示,存储器140或150可存储关联确定模块2202、状态确定模块2204和系统响应模块2206。公开的实施方式并不限于任何特定构造的存储器140。进一步,处理单元110可执行存储在任何模块2202-2206(包括在存储器140或150中)中的指令。

[0276] 在一个实施方式中,关联确定模块2202可存储软件指令(诸如计算机视觉软件),当由处理单元110执行时,该软件指令确定多个交通灯架中的每一个与车辆200的关联。为了本公开的目的,确定交通灯架的关联的步骤可以包括执行一个或更多个评估。在一个实施方式中,处理单元110可以对多个交通灯架中的每一个相对于车辆的定向进行评估。例如,在图21中例示的情况下,图像捕获装置122的视场包括交通灯架2114和交通灯架2122双方。然而,交通灯架2114的定向表明交通灯架2114比交通灯架2122更与车辆200a关联。

[0277] 其它评估可以适用于决定做出处理。例如,处理单元110可以对多个交通灯架中的每一个相对于车辆的距离进行评估。另外地或另选地,处理单元110可以评估哪个交通灯正面向车辆200的前部。另外地或另选地,处理单元110可以使用所识别的车道标记来将车辆200的前方区域划分成多个区域,将每个所识别的交通灯架与区域关联起来,并且评估哪个区域对车辆200最关联。另外地或另选地,处理单元110可以将经由GPS获取的车辆位置与地图数据进行比较,以确定交通灯的关联。例如,车辆200可以对包括关于在多个位置处的可能行驶选项的信息的地图数据进行评估。通过使用GPS获取的车辆位置,处理单元110可以确定哪些行驶选项对接近枢纽的车辆200可用,并且使用该信息来确定在枢纽处的交通灯与车辆200的关联。

[0278] 在一个实施方式中,状态确定模块2204可存储软件指令,当由处理单元110执行时,该软件指令确定包括在被确定为与车辆关联的至少一个交通灯架中的交通灯的状态。

IPR2025-00943
Tesla EX1002 Page 974

如上所述，交通灯的状态还可以与交通灯提供的信息相关联。在典型的交通灯中，发光交通灯的所显示的颜色和相对位置可以提供与车辆200关联的基本信息。

[0279] 状态确定模块2204可以从关联交通灯架的环境中并且从包括在该关联交通灯架中的非典型交通灯中导出附加信息。例如，关联交通灯架可以在其附近具有包括关联文本的标志，例如陈述具体日期和时间的标志。因此，在一些实施方式中，状态确定模块2204可以从包括在与关联交通灯架相关联的标志中的文本中导出信息。例如，状态确定模块2204可以实施光学字符识别技术来识别标志中的文本。然后，状态确定模块2204可以将所识别的文本与数据库进行比较，以确定由标志提供的信息。作为另一示例，关联交通灯架可以包括行人通过信号。状态确定模块2204可以确定行人通过信号的状态是指行人可以在车辆200前面通过道路，而车辆200具有向右转的绿灯。

[0280] 在一个实施方式中，系统响应模块2206可存储软件指令(诸如计算机视觉软件)，当由处理单元110执行时，该软件指令在主车辆200中产生一个或多个响应。一种类型的系统响应可以包括一个或更多个导航响应，如上面参照导航响应模块408详细描述的。其它类型的系统响应可以向节流系统220、制动系统230和/或转向系统240提供控制信号。例如，处理单元110可发送电子信号，致使系统100物理地应用制动器达预定的量或者部分地松开车辆200的加速器。进一步，处理单元110可发送致使系统100将车辆200沿特定方向转向的电子信号。这种响应可以基于关联交通灯的确定。进一步，如果行人可以在车辆200前面通过道路，则系统响应可以包括执行图像数据的附加处理。例如，处理单元110可以确认在车辆200附近的行人在车辆200正在移动时不在将使行人置于危险中的轨迹上行走。

[0281] 枢纽中的交通灯架的位置能够在该枢纽之前、之后或中间。识别枢纽中的每个交通灯架的位置可以被用来例如确定交通灯架的关联。在一些实施方式中，处理单元110可以估计一个或更多个交通灯架相对于车辆200的距离，以创建枢纽的3D模型。在一个实施方式中，枢纽的3D模型可以被存储用于将来使用。枢纽的3D模型可以包括以下各项中的一个或更多个：一个或更多个交通灯架的3D位置、枢纽中的每个交通灯架与其它交通灯架之间的相对距离、每个交通灯架所指示的(多个)方向、以及枢纽的每个交通灯架和停止线之间的相对距离。另外，当车辆200接近枢纽时，可以使用可识别的细节来定期地对3D模型进行更新。可识别细节的示例可以包括交通灯中的箭头、枢纽中的车道标记等等。在一个示例中，当车辆200经过枢纽时，将已识别的细节与存储在3D模型中的信息进行比较，并且如果合适的话，则对3D模型进行更新。

[0282] 与本公开的实施方式相一致，3D模型可以被用来确定一个或更多个交通灯架的关联。为了利用足够的时间将制动系统230激活以使车辆200在枢纽之前停止，系统100可以在距枢纽大约八十米的距离处确定哪个交通灯与车辆200关联。随着车辆200接近枢纽，从所捕获的图像中导出的信息可以与3D模型进行比较以找到匹配。例如，处理单元110可以将枢纽中的所识别交通灯架之间的相对距离与3D模型进行比较，以确定一个或更多个交通灯架的关联。使用3D模型，处理单元110可以在车辆200距枢纽超过50米、75米、100米或125米远时识别关联交通灯。作为另一示例，处理单元110可以将枢纽中的所识别交通灯架之间的相对距离与3D模型进行比较，以确定距枢纽的停止线的距离，即使当该停止线从车辆200的当前位置看不到时。使用3D模型，处理单元110可以在车辆200距枢纽超过50米、75米、100米或125米远时确定距枢纽的停止线的距离。下面将结合图23提供关于关联确定模块2202、状态

IPR2025-00943
Tesla EX1002 Page 975

确定模块2204和系统响应模块2206的进一步细节。

[0283] 图23是示出与公开的实施方式相一致的用于识别关联交通灯的示例性处理2300的流程图。系统100可以实施处理2300来处理在例如图20和图21中描述的情况。

[0284] 如图23中例示的，在步骤2310，图像捕获装置122可以获取车辆200的前方区域的图像。该区域可以包括多个交通灯架，每个交通灯架包括至少一个交通灯。在步骤2320，处理单元110可以例如基于从所获取的图像中识别的特征(例如，车道标记)来确定车辆位置的至少一个指示符。在步骤2330，处理单元110可以使用至少一个指示符来确定多个交通灯架中的每一个的关联。在步骤2340，处理单元110可以确定包括在关联交通灯架中的交通灯的状态。例如，交通灯的状态可以指示车辆200应当停止(例如，由于红灯)或者能够向左转(如果没有正在通行的行人)。在确定包括在交通灯架中的交通灯的状态之后，在步骤2350，处理单元110可以基于所确定的状态变化来产生系统响应。处理2300的这些步骤下面将更加详细地进行讨论。

[0285] 例如，在步骤2310，图像捕获装置122可以获取车辆200的前方区域的至少一个图像。“车辆的前方区域”包括位于车辆的前面的相对于其移动方向的地理区域。区域可以包括枢纽(例如，在图20中为十字路口2000并且在图21中为2100)、十字路口、交叉路、圆形交叉路口、街道、道路等等。在一些情况下，区域可以包括多个交通灯架，每个交通灯架包括至少一个交通灯。“交通灯架”包括容纳用来控制交通和/或提供道路相关信息的一个或更多个光产生装置的任何形式的结构。在一些情况下，两个或更多个交通灯架可以被结合在单个交通灯组件中，但是每个交通灯架可以与不同车道相关联。典型的交通灯架可以包括三个圆形交通灯：绿色交通灯、黄色交通灯和红色交通灯。“交通灯”包括具有能够显示独特颜色的至少一个光源的装置。在一些情况下，车辆200可以遇到非典型交通灯架。非典型交通灯架可以包括具有不同颜色的一个或更多个非圆形交通灯。例如，右转弯箭头交通灯、左转弯箭头交通灯、公共交通交通灯、行人通过交通灯、自行车通过交通灯等。

[0286] 在一些实施方式中，处理单元110可以经由数据接口128接收至少一个所获取的图像。如参照图5D详细讨论的，处理单元110可以识别至少一个所获取图像中的交通灯。例如，处理单元110可过滤至少一个所获取的图像中所识别的对象，以构建一组候选对象，排除不可能对应于交通灯的这些对象。过滤可基于与交通灯关联的各种属性(诸如形状、尺寸、质地、(例如，相对于车辆200的)位置，等)来完成。另外，处理单元110可以分析车辆200的前方区域的几何图形。该分析可基于以下的任何组合：(i)在车辆200的任一侧的检测的车道数目；(ii)在道路上检测的标记(诸如箭头标线)；以及(iii)从地图数据中提取的区域的描述。例如，如果在道路上识别出限定行驶车道的车道标记，并且交通灯架在车道标记的边界之内，则系统100可以得到结论：交通灯架与和车道标记相关联的车道关联。

[0287] 在步骤2320，处理单元110可以从图像中确定车辆位置的至少一个指示符。“车辆位置的指示符”包括与车辆的物理位置相关的任何形式的信息。车辆位置的指示符可以从对图像数据(例如，至少一个所获取图像)的分析中导出。另外，车辆位置的指示符可以从地理位置数据(例如，GPS信号、当地定位信号和/或地图数据)中或者从指示车辆200相对于道路上的其它车辆的位置的数据中导出。在一些实施方式中，车辆位置的指示符可以包括从至少一个图像中导出的距至少一个交通灯架的距离。在其它实施方式中，车辆位置的所述至少一个指示符可以包括基于所述至少一个图像的分析而识别的车道标记。例如，处理单

IPR2025-00943
Tesla EX1002 Page 976

元110可以使用从至少一个图像中导出的信息来进行分析，以识别一个或更多个车道标记。使用所识别的车道标记，处理单元110可以确定所检测的交通灯与车辆200当前正在行驶的车道之间的对应关系。

[0288] 在步骤2330，处理单元110(例如，经由关联确定模块2202)可以使用车辆位置的至少一个指示符，如在步骤2320中所确定的，以确定多个交通灯架中的每一个与车辆200的关联。在一些实施方式中，处理单元110可以对至少一个所获取图像中的所识别交通灯架的关联进行排序。具有关联排序的最高值的交通灯架可以被确定为关联交通灯架。例如，在图20中例示的情况下，图像捕获装置122的视场包括交通灯架2012和交通灯架2014，但是通过使用一个或更多个评估，处理单元110可以确定交通灯架2014具有更高的关联排序值。处理单元110用来确定交通灯架2014具有比交通灯架2012更高的关联排序的一种方法是通过对多个交通灯架中的每一个相对于车辆200a的距离进行评估。例如，在图20中描述的情况下，交通灯架2014更靠近交通灯架2012。因此，交通灯架2014比交通灯架2012更有可能关联。确定关联的其它方法在上文中参照关联确定模块2202进行了描述。

[0289] 在一些实施方式中，多个交通灯架中的每一个与车辆200的关联确定可以包括初步检查，以消除不大可能的交通灯架。例如，当至少一个所获取图像包括三个较近交通灯架和两个较远交通灯架时，两个较远交通灯架可以被归类为不大可能与车辆200关联。通过消除不大可能的交通灯架以及对至少一个所获取图像中的所识别的交通灯架的主体的关联进行排序，处理单元110可以节省处理电力。在一些实施方式中，当车辆200接近枢纽时，关联排序可以发生变化。例如，交通灯架的定向可以从道路上的不同点开始变化，因此，距枢纽的距离可以对给定交通灯架与车辆200关联的概率产生影响。因此，关联排序可以与置信水平相关联，当对交通灯架进行评估时，该置信水平可以考虑诸如距枢纽的距离这样的因素。进一步，当置信水平在特定预定阈值以下时，处理单元110可以定期地或不断地对交通灯架的关联排序进行更新。

[0290] 为了确定与车辆200关联的交通灯，关联确定模块2202还可以使用之前存储在系统1000内部(诸如存储在存储器140/150内部)的导航信息。基于如上讨论的所确定的车辆200的3D位置和/或图像捕获装置122-126，关联确定模块2202可以执行导航地图数据与车辆200之间的注册(registration)。基于如上所述的相对距离测量的确定，关联确定模块2202可以使用导航地图数据来确定车辆200(经由制动系统230)应当进行制动以便在枢纽处的每个所检测的交通灯架处停止的3D空间中的点。根据3D注册结果，关联确定模块2202可以使用导航地图数据来确定在十字路口处检测到的每个交通灯的车道分配。然后，关联确定模块2202可以确定系统100和图像处理器190内部的车道分配，然后执行注册。

[0291] 在步骤2340，处理单元110(例如，经由状态确定模块2204)可以基于至少一个所获取的图像来确定包括在被确定为与车辆200关联的至少一个交通灯架中的交通灯的状态。在一些情况下，关联交通灯架可以包括多个照亮交通灯，并且每个交通灯的状态可以取决于交通灯的类型。在一个实施方式中，交通灯的状态仅指其指示的颜色，例如绿色、黄色和红色。系统100可以使用各种技术来识别交通灯的状态。例如，系统100可以识别包括交通灯在内的一个或更多个图像的区域，并且执行该区域中的像素的分析以确定像素的颜色。在对区域中的至少阈值数量的像素(例如，两个像素、十个像素、二十个像素等等)进行分析之后，系统100可以例如通过发现该区域中的像素的平均值来确定交通灯的颜色。

IPR2025-00943
Tesla EX1002 Page 977

[0292] 另外地或另选地，交通灯的状态可以包括交通灯所指示的方向。在典型交通灯中，绿色指示允许车辆200前进。然而，有时候该信息本身不足以决定在一定方向上(诸如当绿灯仅授权转弯时)行驶是否安全。确定交通灯指示哪个方向的一种方法包括对使每个交通灯和一个或更多个方向相关的数据库进行评估。另一种方法包括从图像数据中识别交通灯的类型以及从环境情境中确定交通灯所指示的方向。例如，在图20中相对于车辆200b描述的第二种情况下，交通灯架2024可以被确定为关联交通灯架，但是交通灯架2024包括两个照亮的交通灯2024a和2024b。因此，在步骤2340中，处理单元110可以确定交通灯2024a的状态是用于继续直行的红灯，并且交通灯2024b的状态是用于右转弯的绿灯。在一些实施方式中，交通灯的状态的确定包括以下步骤中的一个或多个：确定在关联交通灯之内的交通灯的位置；确定交通灯是否被照亮；确定交通灯的颜色；以及确定交通灯是否包括箭头。

[0293] 在步骤2350，处理单元110(例如，经由系统响应模块2206)可以基于所确定的状态来产生系统响应。在一些实施方式中，处理单元110可以产生一个或更多个系统响应(例如，两个或更多个响应)，包括不同类型的响应。一种类型的系统响应可以包括导航响应，该导航响应可以包括例如开始行驶、加速度的变化、速度的变化、施加车辆制动、中止巡航控制等等。例如，这些系统响应可以包括向节流系统220、制动系统230和转向系统240中的一个或更多个提供控制信号以对车辆200进行导航。另一种类型的系统响应可以包括启动定时器和/或增加计数器，以便向车辆200的驾驶员提供关于一个或更多个驾驶过程的统计信息。例如，统计信息可以指示在驾驶过程中车辆200已经遇到红灯的次数，和/或车辆200在红灯处等待花费的驾驶过程的持续时间。另一种类型的系统响应可以包括向车辆200的驾驶员提供各种通知(例如，警告和/或警报)。警告和/或警报可以包括例如通知关联交通灯的颜色和/或距枢纽的距离。通知可以经由扬声器或者经由关联显示器(例如，触摸屏320)来提供。

[0294] 转弯车道交通灯检测

[0295] 系统100可以检测车辆(诸如车辆200)在转弯车道上行驶。车道可指车辆的指定或预期行进路径并且可具有标记边界(例如，道路上的线)或未标记边界(例如，道路路边、道路障碍、护栏、停放的车辆，等等)或约束。例如，系统100可以基于经由一个或更多个图像捕获装置获取的视觉信息来进行操作以做出这些检测和确定。在一些实施方式中，还可以进行至少部分地关于地图数据和/或所感测的车辆位置的这些检测和确定。除了确定在转弯车道中的车辆200的状态以外，系统100可以识别与车道相关联的交通灯，并且可以被配置为基于道路环境信息的分析和所确定的交通灯的特性来确定交通灯的状态。

[0296] 图24例示了在道路800上行驶的车辆200，其中可以使用用于检测道路上的对象的所公开的系统和方法。车辆200被描述为配备有图像捕获装置122和124，可以在任何特定的车辆200上采用更多或更少相机。如所示，道路800可以被细分成车道，诸如车道810和820。作为示例示出了车道810和820；给定道路800可以基于道路的大小和性质而具有额外车道。在图24的示例中，车辆200在车道820上行驶，并且能够看到其正在接近与交叉路2404的十字路口2402。在十字路口2402处车道820上的交通通过交通灯架2406来控制。如下面详细讨论的，系统100可以确定交通灯架2406的状态并且产生影响车辆200的操作的系统响应。道路800上进一步例示的是通向停止线2412的警告线2410。

[0297] 处理单元110可以被配置为基于处理单元110可经由数据接口128接收的由一个或

IPR2025-00943
Tesla EX1002 Page 978

更多个图像捕获装置122-126获取的多个图像来确定与车道810和820、十字路口2402或者交义路2404中的每一个相关联的一个或更多个车道约束。根据一些实施方式，车道约束可由可见的车道边界(诸如标记在路面上的虚线或实线)识别。另外地或替代地，车道约束可包括路面的路边或障碍。另外地或替代地，车道约束可包括标记(例如，博茨点)。根据一些实施方式，处理单元110(经由下面详细描述的车道定位模块2520)可以通过识别诸如道路800或交叉路2404的整体的路面宽度的中点或者车道810/820中的一个的中点来确定与车道810/820、十字路口2402或交义路2404相关联的约束。当例如没有画出或者标记出指定诸如车道810/820的道路车道的线时，处理单元110可以基于已知的道路参数来通过替代的方式(诸如通过估计或推断)来识别车道约束。处理单元110还可以确定各种约束和所检测的对象(诸如交通灯)之间的物理相对距离。

[0298] 当建立自主车辆或红色交通灯警告系统时，对于具有交通灯的枢纽的距离估计可能具有挑战性，这是因为交通灯的位置可以在该枢纽之后或者中间。例如，确定针对十字路口处的每个交通灯的3D位置可以暗示对于车辆的制动点，但是停止线可以提供精确位置。

[0299] 在当前的系统中，道路800、交叉路2404或十字路口2402以及构成车道810/820的约束和路径的检测可以包括经由相机坐标系来确定其3D模型的处理单元110。例如，车道810/820的3D模型可由三次多项式描述。除了行驶车道的3D建模以外，处理单元110还可执行主机运动参数(诸如车辆200的速度、偏航和俯仰速率以及加速度)的多帧估算。处理单元110还可确定道路标高模型以将从多个图像获取的信息转换为三维空间。当前系统的一个特征在于，可以同时创建针对交通十字路口处的静止对象的总体模型，而不是创建针对每个所检测对象的一个单独模型。因此，系统100可以被配置为确定针对在例如在大约一百米之内的距离处的对象或车道约束的绝对距离。在一些情况下，系统100可以被配置为确定针对其它距离(例如，在125米内、150米内等等)处的对象或车道约束的绝对距离。进一步，系统100可以被配置为在距停止线2412大约八十米的距离之内确定哪个交通灯架(例如，交通灯架2406)与车辆200关联。

[0300] 可以通过被称为扩展和缩放的处理来实现对十字路口处的交通灯的精确距离估计。扩展处理可以使用相对距离来确定从车辆200到对象(诸如停止线2412)的距离。例如，如果在图像数据中检测到两个交通灯架，并且当车辆移动十米的距离时两个交通灯架之间的距离增加5%，则系统计算出距离为200米。

[0301] 图25是与公开的实施方式相一致的存储器140和/或150的示例性框图，该存储器140和/或150可存储用于执行一个或多个操作的指令。如图25中例示的，存储器140/150可存储用于执行本文中描述的对象检测和响应的一个或多个模块。例如，存储器140/150可以存储交通灯检测模块2510、车道定位模块2520和动作模块2530。应用处理器180和/或图像处理器190可执行存储在任何模块2510-2530(包括在存储器140/150中)中的指令。本领域技术人员将理解的是，以下讨论对处理单元110的引用可单独或共同指应用处理器180和图像处理器190。因此，任何以下处理的步骤可由一个或多个处理装置执行。

[0302] 交通灯检测模块2510可以存储指令，当由处理单元110执行时，该指令可以检测交通灯架(诸如交通灯架2406)的存在和状态。如下面将结合图26A、图26B和图27进行讨论的，交通灯检测模块2510连同图像处理器190可以执行对由一个或更多个图像捕获装置122-126获取的一个或更多个图像的图像处理。交通灯检测模块2510还可以确定交通灯架2406

IPR2025-00943
Tesla EX1002 Page 979

的状态，包括确定交通灯架2406内部的一个灯是否包括箭头。交通灯检测模块2510可以确定与交通灯架2406的状态关联的其它信息，非限制性地包括交通灯架2406内部关联的任何交通灯是否被照亮(以固定或闪烁的方式)、确定交通灯架内部的交通灯的位置(即，交通灯的水平定向与垂直定向)、或者确定与交通灯相关联的颜色。在一些实施方式中，交通灯检测模块2510可以存储被确定用于处理单元110内部(诸如存储器140/150中)的特定前进方向、特定十字路口2402和/或特定交通灯架2406的信息。在这些实施方式中，当车辆200返回至同一十字路口时，之前确定且保存的信息可以在未来使用。

[0303] 在一些实施方式中，可以通过在预定或已知捕获率获取的多个图像的图像分析(例如，对以1秒、1.5秒、2秒等等获取的图像分别进行分析)来确定闪烁的交通灯。例如，系统100可以分析图像数据，以识别多个图像当中的发光图案。系统100还可以确定被确定为在交通灯架中的受关注的特定交通灯的边界之内的所捕获的图像的区域。然后，系统100可以通过被确定为在受关注的交通灯的边界之内的区域中的像素分析来确定交通灯的颜色。

[0304] 车道定位模块2520可存储指令，当由处理单元110执行时，该指令可帮助系统100确定车辆200的位置。在一些实施方式中，确定车辆200的位置的步骤可以包括：经由视觉确定或者通过经由数据接口128从图像捕获装置122-126接收的至少一个图像的分析来确定车辆位置的至少一个指示符。在这些实施方式中，车辆位置的至少一个指示符可以包括从车辆到与车辆正在行驶的当前车道相关联的一个或更多个车道约束或车道标记，或者到与十字路口相关联的标记(诸如警告线2410和停止线2412)的距离。从车辆到位置或对象的距离可以基于例如图像数据、GPS信息或来自位置传感器的数据的分析中的一种或多种来确定。进一步，车辆位置的至少一个指示符可以包括与具有至少一个关联交通灯的交通灯架相关联的箭头。另外地或另选地，车辆位置的至少一个指示符可以包括通过GPS或者类似坐标系获取的车辆定位。在一些实施方式中，车道定位模块2520可以存储针对特定道路800及其车道810/820而确定的车道约束信息。例如，处理单元110可以将信息存储在存储器140/150中。在这些实施方式中，当车辆200返回至同一十字路口时，之前确定且保存的信息可以在未来使用。例如，GPS信息可以被用来确定车辆200已经回到同一十字路口。

[0305] 与公开的实施方式相一致，车道定位模块2520可以使用来自车辆位置的至少一个指示符来确定是否需要或推荐改变车辆200的操作的系统响应。另外地或者另选地，车道定位模块2520可以从其它模块(包括位置传感器130)或其它系统接收指示存在图像数据中的其它特征的信息(诸如其它车辆、道路的曲率等)。

[0306] 动作模块2530可以存储指令，当由处理单元110执行时，该指令可以帮助系统100基于从一个或更多个源(诸如位置传感器130、图像处理器190、交通灯检测模块2510)接收的信息来进行相对于车辆200的操作的一个或更多个动作。在一些实施方式中，动作模块2530可以(从例如交通灯检测模块2510)接收关于十字路口处的交通灯(诸如如上所讨论的十字路口2402处的交通灯架2406)的状态的信息。由动作模块2530接收的其它信息可以包括车辆是否在转弯车道以及交通灯架2406是否包括箭头的确定。

[0307] 基于该所接收的信息，动作模块2530然后可以产生影响车辆200的操作状态的系统响应，诸如促使系统100向节流系统220、制动系统230和转向系统240中的一个或多个提供控制信号，以(例如，通过加速、转弯，等等)导航车辆200。例如，在一些实施方式中，系统100可以确定转弯车道交通灯授权车辆200进行转弯。在这些实施方式中，动作模块2530可

IPR2025-00943
Tesla EX1002 Page 980

以向转向系统240发送指令，并且转向系统240可以执行该指令以使车辆200转弯通过十字路口2402进入与交叉路2404相关联的新行驶车道。在其它实施方式中，由动作模块2530发起并执行的系统响应可以包括以下各项的任何或全部：向车辆的操作者提供视觉或听觉通知、施加车辆制动、中止巡航控制功能或者发起一个或更多个自动转弯机动。

[0308] 另外地或者另选地，动作模块300可以向与车辆200相关联的其它系统（诸如制动系统230、转弯信号、节流系统220等）发送指令。在一些实施方式中，动作模块2530相反可以为车辆的人类操作者提供代表从关联系统和/或传感器收集的信息的音频、视觉、或触觉反馈。然后，人类操作者可以对该反馈进行动作，以使车辆转弯。

[0309] 图26A提供了图24中描述的情况的带注释的视图。车辆200再次在道路800的车道820上行驶，并且正接近与交叉路2404的十字路口2402。车辆200再次配备有图像捕获装置122和124；更多或更少装置可以与任何特定车辆200相关联。为了简单进行例示，图26A中的道路800被描述为具有两个行驶车道810和820的在页面上从上到下定向的单向街道。交叉路2404是具有在每个方向上前进的一个行驶车道的穿过整个页面的左右定向的双向街道。

[0310] 在十字路口2402处向道路800的左侧的是交通灯架2406。在车辆200的前面的车道820的表面上画出的车道箭头2408，指示车道820是左转弯车道。同样在道路800的表面上画出或者附加的是通向停止线2412的警告线2410。

[0311] 与公开的实施方式相一致，系统100可以被配置为确定接近十字路口的车辆200的行驶车道（这里，车道820）是否是转弯车道；确定交通灯架（这里，交通灯架2406）是否对十字路口做出控制；确定交通灯架中的交通灯的状态；以及确定交通灯是否包括箭头。下面将结合图27和处理2700来更加详细地对该处理进行描述。

[0312] 简单来说，与车辆200相关联的系统100可以经由处理单元110、位置传感器130、交通灯检测模块2510或车道定位模块2520中的一个或更多个来确定道路800内部的车辆200的位置。另外地或另选地，系统100可以收集来自车辆位置的至少一个指示符的信息。如上面结合车道定位模块2520所讨论的，车辆位置的至少一个指示符可以包括从车辆200到与车辆正在行驶的当前车道（诸如车道820）相关联的一个或更多个车道约束或车道标记，或者到与十字路口2401相关联的标记（诸如警告线2410和停止线2412）的距离。进一步，车辆位置的至少一个指示符可以包括与交通灯架2406相关联的箭头。另外地或另选地，车辆位置的至少一个指示符可以包括通过GPS或者类似坐标系获取的车辆定位。

[0313] 在图26A的例示中，存在车辆200的位置的多个指示符。与车辆200相关联的一个或更多个图像捕获装置122-126可以被配置为捕获可以帮助确定车辆200的位置的在车辆200的前面的区域的多个图像。例如，图像可以包括车道箭头2408，其指示车道820是十字路口2402的左转弯车道。图像还可以包括交通灯架2406，并且可以指示交通灯架2406内部的各个灯中的一个或多个包括表明转弯车道状况的箭头。一个或更多个图像捕获装置122-126还可以捕获与道路800或者交叉路2404相关联的警告线2410或停止线2412的图像。进一步，系统100经由车道定位模块2520可以通过测量从车辆200到停止线2412的距离s_1来确定转弯车道内部的车辆200的位置。仍进一步，系统100可以实施光学字符识别（OCR）处理以获得包括在一个或更多个所捕获的图像中的文本（例如，来自标志和/或道路标记的文本）。然后，系统100可以使用文本信息作为确定车辆200是否在转弯车道内部的一部分或者作为其基础。例如，系统100可以识别指示转弯车道的特定词语（例如，“向右转弯”、“向左转弯”等

IPR2025-00943
Tesla EX1002 Page 981

等)。

[0314] 测量可以是直接测量(诸如经由位置传感器130),或者可以通过图像处理器190经由所捕获的图像数据的分析来确定,并且可以与地图数据相结合地使用。测量结果可以可以从内部或外部车辆200的任何部分测量出,该任何部分非限制性地包括:车辆200的前部、车辆200的诸如前灯或前面车牌的部分、图像捕获装置122-126的安装好的位置、所确定的车辆200的形心、车辆200的后部、与车辆200相关联的一个或更多个挡风玻璃或反光镜、车辆200的车轮、车辆200的右侧或左侧或车窗、与车辆200的车顶相关联的点、或者与车辆200的底盘相关联的点。

[0315] 在一些实施方式中,确定从车辆200到交通灯架2406的距离可能足以帮助处理单元110计算制动距离或者其它这种测量。然而,在其它实施方式中,交通灯架2406可以被定位为经过车辆200将被要求停止的点。在这些实施方式中,可以通过使用警告线2410或停止线2412中的一个或更多个来确定距十字路口的距离。

[0316] 如上面结合图24所讨论的,可以使用采用扩展和缩放的估计处理来计算这些距离。在具有两个交通灯架的十字路口2402处,例如,系统100可以对如在图像数据中捕获的两个架之间的相对距离进行一定时间段的测量,然后使用该相对距离来估计距停止线2412的距离。在一些实施方式中,这些测量可以随着时间的推移重复,以提高精确度。精确度还可以通过其它方法来提高:例如,如果在十字路口2402处坐落有三个或四个交通灯架,则可以计算架中的每一个之间的相对距离并取平均值。另外地或另选地,如上所讨论的,可以使用卡尔曼过滤器来估计距离。

[0317] 系统100仍然需要附加信息或输入来确定停止线2412相对于十字路口2402和/或交通灯架2406的定位。在一些实施方式中,系统100可以使用从之前到十字路口的路程中收集的之前存储的地图数据。该信息可以从交通灯检测模块2510或者车道定位模块2520中接收。还可以使用下式从图像数据中导出距离测量值Z:

[0318] $Z = fW/w$

[0319] 其中,W是两个交通灯架2406之间的已知距离,w是像素中的图像数据中的距离,并且f是像素中的特定图像捕获装置122-126的焦距。

[0320] 在图26A中,插图2602示出与示例关联的交通灯架2406的两种状态。基于经由数据接口128从图像捕获装置122-126接收的来自至少一个图像的图像数据,系统100可以被配置为确定交通灯架2406的状态。下面将结合图27和处理2700来更加详细地对该确定进行描述。当接近十字路口2402时,系统100可以首先确定交通灯架2406止显示固定的红灯,如在插图2602的图像1中所见。在图26A的示例中,车道820是左转弯车道,所以除非道路800位于红灯时左转弯是合法的管辖范围内(或者如果交叉路2404被配置作为在页面上具有从右向左流动的两个车道的单向街道的话),否则车辆200必须在停止线2412处停止并且保持静止直至确定交通灯架2406已经发生改变。在另选实施方式中(未示出),车辆200可以在被确定为右转弯车道的车道上行驶,并且假设红灯时右转弯在该管辖范围内是合法允许的,则系统100可以产生系统响应,使得车辆200能够右转弯同时交通灯架2406保持红色。

[0321] 随后,如从插图2602的图像2中所见,交通灯架2406从固定红灯转变成指示受保护的左转弯的固定绿色箭头。系统100可以被配置为基于通过图像处理器190对交通灯架2406的所获取图像的像素级图像分析来不但检测出灯的颜色已经改变,而且检测出箭头已经出

IPR2025-00943
Tesla EX1002 Page 982

现在灯内部。在一些实施方式中，为了使系统100检测出交通灯的颜色已经改变，系统100可以对所捕获的图像进行一定时间段(例如，1秒、5秒、10秒、15秒、20秒、30秒、1分钟等等)的分析，以便识别交通灯的状态变化。系统100还可以一次确定交通灯架2406发生变化以显示其是否安全或者是否被授权在由灯的状态指示的方向上前进的绿色箭头。即使灯可以显示绿色箭头，各种情况(未示出)也可能要求车辆200在完成转弯之前进一步延迟。例如，行人可能在通过交叉路2404。虽然在图26A中道路800是单向街道，但是在另选实施方式中(诸如下面描述的且在图26B中示出的)，道路可以是双向街道，并且车辆200还可以对从相反方向驶来的第二车辆是否正在从另一方向在红灯时进行右转弯到交叉路2404上来。进一步，其它障碍物可以阻止及时的左转弯；例如，一组铁路轨道可以与道路800平行地向其左侧定向，并且虽然左转弯绿色箭头被照亮，但是与火车轨道相关联的障碍可能落下来以允许火车通过。这些示例不用于限制。一旦系统100确定了交通灯架2406和十字路口处的状况都允许安全转弯，系统100就可以经由动作模块2530产生涉及节流系统220、制动系统230、转向系统240、转弯信号、和/或使车辆200自动向左转弯的其它子系统和模块中的一个或更多个的系统响应。

[0322] 图26B例示了与公开的实施方式相一致的其它十字路口场景。在图26B的示例中，车辆200(这里表示为车辆200e)正在道路800的车道810上行驶，并且接近与交叉路2404的十字路口2402。车辆200再次配备有图像捕获装置122和124；更多或更少装置可以与任何特定车辆200相关联。与图26A的例示不同，道路800现在在图26B中是双向街道，并且交通可以在两个方向上进行。交叉路2404依然是具有在每个方向上前进的一个行驶车道的穿过整个页面的左右定向的双向街道。

[0323] 在十字路口2402处向道路800的右侧的是交通灯架2604。在车辆200的前面的车道810的表面上画出的车道箭头2408，指示车道820是左转弯车道。如前，与公开的实施方式相一致，系统100可以被配置为确定接近十字路口的车辆200的行驶车道(这次，车道810)是否是转弯车道；确定交通灯架(这里，交通灯架2604)是否对十字路口做出控制；确定交通灯的状态；以及确定交通灯是否包括箭头。

[0324] 系统100可以按照与上面描述的大致相同的方式来进行操作，以进行这些确定。然而，在图26B的示例中，插图2606证明当车辆200e接近十字路口2402时，交通灯架2604正显示闪光或闪烁的黄灯。在插图2606中通过箭头的虚线指示非固定性质。如可能对相关技术领域的普通技术人员来说熟悉的，一些管辖范围在一天的某些时间期间使用闪烁的黄灯来实质上用作与“让行”或“停止”标志相同的功能，同时仍然允许利用红灯使行驶车道完全停止或者利用绿灯完全授权行驶车道前进。

[0325] 如上面概述的，系统100可以确定交通灯架2604正在显示闪烁的黄灯，并且还可以确定车道810是转弯车道。然而，在图26B中描述的情况下，系统100可以被要求进行其它确定，以便使车辆200e成功穿过十字路口2402。如本领域技术人员将理解的，闪烁的黄灯授权接近的车辆随意转弯，只要这样做安全的话。希望驾驶员在十字路口处要谨慎，并且驾驶员确定太早进行转弯的话将不会产生潜在的碰撞或其它危险。

[0326] 系统100可以进行的一个这种确定是一个或更多个车辆是否正在从另一方向接近十字路口。例如，尽管在道路的一个方向上行驶的行驶车道可以具有转弯车道中的闪烁的黄灯的交通灯，在相反方向上行驶的交通可以具有绿灯以直行前进通过十字路口。因此，如

IPR2025-00943
Tesla EX1002 Page 983

果车辆的操作者在闪烁的黄灯信号下没有让行的话，碰撞可能发生。在图26B的示例中，第二车辆200f确实正在直行通过十字路口2402。系统100可以被配置为经由从一个或更多个图像捕获装置122-126捕获的图像数据来检测车辆200f的存在，以及使用位置传感器130、图像处理器190、交通灯检测模块2510或车道定位模块2520中的一个或更多个来确定车辆在道路/十字路口内的位置以及其速度。在检测并分析车辆200f时，系统100(经由动作模块2530)可以促使制动系统230保持接合直至车辆200f已经通过。

[0327] 一旦没有其它车辆正在从相反方向接近十字路口，系统100就可以经由动作模块2530基于该确定来产生系统响应。系统响应可以包括以下中的一种或多种：向车辆200的操作者提供反馈、经由制动系统230来施加车辆制动、中止巡航控制或者使用节流系统220、制动系统230或转向系统240中的一个或更多个来发起至交叉路2404上的自动转弯机动。首先，然而，系统100可以分析针对与十字路口2402或交叉路2404相关联的任何其它障碍物或对象的存在的图像数据。例如，在图26B的例示中，能够看到行人2608步行穿过交叉路2404，因为道路800的另一个车道具有该方向上的绿灯。系统100可以通过分析由一个或更多个图像捕获装置122-126捕获的多个图像中的一个或更多个来经由图像处理器190检测行人2608的存在。如上所述，系统100的其它子系统和模块可以配合以检测行人2608的速度，并且可以通知系统100行人2608何时不再在道路上并且转弯是安全的。

[0328] 图27例示了与公开的实施方式相一致的用于检测交通灯的状态的处理2700。处理2700的步骤可以通过处理单元110、图像获取单元120、位置传感器130、图像处理器190、交通灯检测模块2510、车道定位模块2520或动作模块2530中的一个或更多个来执行。

[0329] 在步骤2710，处理2700可包括：使用至少一个图像捕获装置122、124和/或126来获取车辆200的前方区域的多个图像。该区域可以包含具有至少一个交通灯的一个或更多个交通灯架(诸如上述的交通灯架2406和2604)。处理单元110可以通过数据接口128从(多个)图像捕获装置接收多个图像。然后，多个图像可以通过处理单元110的图像处理器190来实时地进行处理。图像处理器190可以对由至少一个图像捕获装置捕获的至少一个图像、以及在稍后的时间点捕获的其它图像进行处理。

[0330] 在步骤2720，处理2700可以基于车辆位置的至少一个指示符来确定车辆200是否正在转弯车道上行驶。如上面通常描述的，可以使用位置的各种指示符来确定行驶车道是否是转弯车道。在一些实施方式中，位置指示可以包括通过GPS获取的车辆定位。在其它实施方式中，系统100可以识别车道标记或约束(诸如指示车道是转弯车道的车道箭头2408)。另选地，图像处理器190可以确定箭头与控制车道行驶的交通灯架相关联。另外，可以使用系统100的各种元件来进行距离的测量。例如，位置传感器130和/或车道定位模块2520可以将包括车辆200距十字路口2402有多远(上述的距离s_1)在内的关于车道810/810内部的车辆200的相对位置的信息发送到处理单元110。车道的尽头可以基于所测量的距离或者视觉确定(例如，通过对由一个或更多个图像捕获装置122-126获取的一个或更多个图像进行分析以及检测警告线2410或停止线2412的存在)。

[0331] 在步骤2730，处理2700可以执行对经由数据接口128接收的至少一个图像的图像处理，以确定箭头是否存在于与交通灯相关联的图像数据中。交通灯检测模块2510和/或图像处理器190可以被配置为分析与代表交通灯架内部的各个交通灯的位置相对应的图像的像素。一种或更多种图像处理算法可以被应用于像素，以确定固定或闪烁的/闪光的箭头是

IPR2025-00943
Tesla EX1002 Page 984

否存在于关联区域中。下面将结合图28至图30B来讨论系统100可以怎样对一个或更多个图像进行分析以确定交通灯是否显示箭头的进一步示例。

[0332] 在步骤2740，处理2700可以确定交通灯的状态。如上所讨论的，确定交通灯的状态的步骤可以非限制性地包括：确定交通灯是否被照亮(和/或是否起作用)、确定交通灯架内部的被照亮交通灯的位置、或者确定与交通灯相关联的颜色(即，红色、黄色或绿色)。还列表仅用于示例，并且可以通过系统100和处理单元110来确定并处理与交通灯关联的任何信息。例如，系统可以被配置为基于对一系列获取图像的图像处理来确定交通灯是固定的还是闪烁的。图像处理器190可以确定在这些实施方式中，架内部的同一交通灯在一个图像中被照亮，按顺序接下来变暗，然后即刻被再次照亮；这将指示特定的灯正在闪烁。在一些实施方式中，系统100可以被配置为识别对车辆正在行驶的特定管辖范围来说特别的交通灯状况。该信息可以被预先编程到系统100中，诸如存储在存储器140/150内部的数据(未示出)中。

[0333] 在步骤2750，处理2700可以经由对车辆200的操作产生影响的系统100来产生系统响应。由系统100发起的响应可以基于以下确定中的任何或全部：确定交通灯的状态、确定交通灯是否包括箭头以及确定车辆是否在转弯车道。如上所讨论的，由系统100经由动作模块2530发起的特定动作可以非限制性地包括：向车辆200的操作者提供关于交通灯的视觉或听觉反馈通知；经由制动系统230施加车辆制动；中止巡航控制功能；或者使用节流系统220、制动系统230或转向系统240中的一个或更多个来继续前进以发起自动转弯机动并且继续前进通过十字路口。

[0334] 另外，在一些实施方式中，作为系统100的功能的一部分，工作模块2530和/或处理单元110还可以被配置为确定车辆所处的转弯车道是否被授权在通过所确定的交通灯的状态指示的方向上前进。换句话说，即使交通灯指示特定的系统响应是合法允许的，系统100也可以通过进一步确定这种动作实际上是否可能且安全来为车辆200的驾驶员提供额外等级的安全保护。例如，如上面结合图26B所讨论的，系统100和图像处理器190可以被配置为检测与道路或十字路口相关联的其它障碍物的存在(诸如存在于街道上的行人或者其它车辆的存在)。在一些实施方式中，其它车辆可以正常前进通过十字路口(诸如在图26B中描述的车辆200e具有闪烁的黄灯的场景下)。在这些实施方式中，车辆200e理论上能够转弯到交叉路2404上，但是系统100还可以考虑车辆200f在绿灯时正在前进通过十字路口的事实。另选地，第二车辆可以不寻常或非法的方式出现，诸如如果另一车辆的操作者闯过停车标志或者红灯，或者如果车辆在十字路口或者其中一个道路上已经故障。

[0335] 仍然在其它实施方式中，确定车辆所处的转弯车道是否被授权在由交通灯的所确定的状态指示的方向上前进的步骤还可以包括：基于车辆的预定目的地来确定实际上是否遵守灯的指示状态。在涉及自主车辆的实施方式中，可以假设操作者和/或乘客已经向系统100的处理单元110指示了其期望目的地。因此，通过使用GPS定位技术或其它这种信息，系统100可以被配置为确定进行由交通灯指示的转弯实际上是否与预定目的地关联。例如，在一些实施方式中，车辆200已经无意地在转弯车道上熄火，并且向预定目的地前进可以涉及并入相邻车道(在车道定位模块2520的帮助下)以及继续前进直行通过十字路口。因此，系统100可以被配置为基于该预定目的地信息来确定是向左转弯、向右转弯还是继续直行通过十字路口。

IPR2025-00943
Tesla EX1002 Page 985

[0336] 交通灯细节检测

[0337] 在一些实施方式中，系统100可以提供与交通信号的识别相关的附加功能。例如，系统100可以基于交通灯的物理特征的分析来提供交通灯细节检测功能以提供反馈。为了实现这种功能，处理单元110可以处理由图像捕获装置122、124、126中的至少一个捕获的图像并且采用超分辨率技术来识别图像中的交通灯的特定特征，其中该特定特征可以向系统100指示有用信息。例如，系统100可以识别交通灯的已经被照亮的圆形信号和箭头形信号，并且在其之间进行区分。

[0338] 在一些实施方式中，存储器140和/或150可以存储指令，这些指令被编程为使得当由处理装置执行时将提供交通灯细节检测功能。如图28中所示，存储器140和/或150可以存储图像对齐模块2810、像素扩展模块2820、图像分析模块2830和数据库2840。图像对齐模块2810可以存储用于将由图像捕获装置122、124、126捕获的多个图像进行对齐的指令。像素扩展模块2820可以存储用于增加图像的至少一部分的分辨率以允许图像被缩放成标准尺寸的指令。图像分析模块2830可以存储用于对图像进行分析以检测并识别图像的特征的指令。数据库2830可以被配置为存储与提供交通灯细节检测功能相关联的数据并且当需要时提供数据。进一步，图像对齐模块2810、像素扩展模块2820和图像分析模块2830可以单独地或者以与彼此的各种组合的方式来存储可以由一个或更多个处理器(例如，处理单元110)执行的指令。例如，图像对齐模块2810、像素扩展模块2820和图像分析模块2830可以被配置为与彼此和/或系统100的其它模块相互作用，以执行与公开的实施方式相一致的功能。

[0339] 数据库2840可以包括存储信息并且通过计算装置(诸如处理单元110)进行访问和/或管理的一个或更多个存储装置。在一些实施方式中，如图28中所示，数据库2840可以位于存储器150中。在其它实施方式中，数据库2840可以被定位得距存储器150较远，并且经由一个或更多个无线连接(例如，无线网络)对于系统100的其它部件(例如，处理单元120)可访问。虽然示出了一个数据库2840，但是应当理解，多个单独的和/或互连的数据库可以组成数据库2840。数据库630可以包括计算部件(例如，数据库管理系统、数据库服务器等等)，该计算部件被配置为接收并处理针对存储在与数据库2840相关联的存储器装置中的数据的请求并且(例如，向处理单元110)提供来自数据库2840的数据。

[0340] 在一些实施方式中，数据库2840可以被配置为存储与提供交通灯细节检测功能相关联的数据。例如，数据库2840可以存储数据(诸如图像、地图、算法、多组值等等)，这可以允许处理单元110识别在图像中检测的信息。例如，数据库2840可以存储与交通灯的圆形信号和箭头形信号中的每一个相对应的一组平均像素值。

[0341] 在一些实施方式中，系统100可以提供基于在一个或更多个所捕获的图像中检测的交通灯的细节的反馈。例如，系统100可以被配置为在与交通灯相关联的信号的颜色、形状和位置之间进行区分，以便提供基于所检测的细节的含义的特定反馈(例如，箭头形信号控制转弯移动，而圆形信号控制直行向前移动)。

[0342] 图29A和图29B描述了可以由系统100检测的示例性对象。图29A描述了包括圆形信号2910的交通灯架2900A。图29B描述了包括箭头形信号2920的交通灯架。虽然两个对象都包括交通灯，但是它们包括指示不同消息的不同细节。在示例性实施方式中，系统100可以被配置为对一个或更多个图像的像素进行分析以识别区别特征，诸如圆形信号2910和箭头形信号2920。

IPR2025-00943
Tesla EX1002 Page 986

[0343] 图29C描述了圆形信号2910的图像2915的示例性部分，包括多个像素2930。系统100可以被配置为对像素2930进行分析以识别图像中发现的对象和/或细节。例如，图像分析模块2830可以被配置为识别与每个像素2930相关联的特征(诸如亮度、颜色、对比度、色调、透明度等等)，并且基于像素分析来确定与图像相关联的对象和/或细节。图像分析模块2830识别的每个特征可以被给定一个值。例如，图像分析模块2830可以将值“一”分配给被识别为颜色红色的像素，并且将值“零”分配给任何其它像素。

[0344] 在一些实施方式中，系统100可以被配置为使用超分辨率技术来允许多个图像中的对象的更精确的分析。例如，像素扩展模块2820可以修改图像2915以将像素2930中的一个或更多个划分成子集的像素2940。虽然在图29C中示出仅部分像素2930被扩展到像素2940中，但是应当理解，像素2930的任何部分或全部都可以按照这种方式来划分。因此，像素扩展模块2820可以增加与图像的特定部分相关联的像素的数量。

[0345] 在示例性实施方式中，像素扩展模块2820可以增加与图像中的特定对象相关联的像素的数量，使得图像的包括对象的部分包括所选定数量的像素(例如，标准数量的像素)。通过这种方式，即使多个图像包括不同尺寸的同一对象(例如，因为图像是在距对象不同的距离处拍摄的)，图像也能够被修改为使得图像的包括该对象的部分包括相同数量的像素，该相同数量的像素能够直接进行比较并被平均化，以更加可靠地确定对象的位于像素内部的部分的特征(例如，颜色)。

[0346] 将图像的已选部分缩放以使相同数量的像素与每个图像中的对象相关联可以允许对图像直接进行比较(例如，将图像中一个像素的值与另一图像中相同像素进行比较)。通过这种方式，可以通过图像分析模块2830对多个图像进行考虑和分析，从而与一个图像的分析相比增加准确性。例如，超分辨率技术可以增加图像分析模块2830可以正确地在图像中的圆形信号2920和箭头形信号2930之间进行区分的可能性。

[0347] 图29D描述了箭头形信号2920的图像2925的示例性部分。如上面相对于图像2915所描述的，图像2925可以被划分成像素2930，每个像素2930可以被扩展到像素2940的子集中。如图29C和图29D中能够看出的，像素2930围绕各个信号2920、2930的边缘而不同。图像分析模块2830可以被配置为识别像素2930的一个或更多个特性，以确定像素是对应于圆形信号2910还是箭头形信号2920。像素2930扩展到像素2940中可以允许将考虑多个图像并将多个图像对齐，从而增加检测的可靠性。

[0348] 图30A是示出与公开的实施方式相一致的用于确定交通灯是否包括箭头的示例性处理3000A。在示例性实施方式中，系统100可以执行处理3000A以确定在多个图像中检测的交通灯的部分是否包括箭头形信号。箭头形信号(或者“箭头”)可以是信号的视觉成分采用箭头形状的交通灯的信号的类型。箭头可以是指示此时可以允许车辆转弯(例如，在箭头的方向上转弯)的信号。

[0349] 在步骤3002，处理单元110可以接收交通灯的多个图像。例如，一个或更多个图像捕获装置122、124、126可以获取车辆200的前方区域的多个图像。在一些实施方式中，诸如当车辆200正在接近枢纽时，车辆200的前方区域可以包括包含至少一个交通灯的一个或更多个交通灯架。在这些实施方式中，一个或更多个图像捕获装置122、124、126可以获取交通灯的多个图像。处理单元110可以通过数据接口128从图像捕获装置122、124、126接收多个图像。

IPR2025-00943
Tesla EX1002 Page 987

[0350] 处理单元110可以分析所接收的多个图像。在一些实施方式中，处理单元110可以对多个图像进行分析以确定多个图像包括交通灯。例如，处理单元110可以识别一个或更多个对象作为交通灯的候选，并且执行一个或更多个处理(例如，图像匹配)以确定该候选是否对应于交通灯。然而，应当理解，处理单元110可以执行处理3000A，无需识别对象作为交通灯的不同步骤。另外，处理单元110可以对多个图像中的一部分进行分析以确定交通灯的特性。这些特性可以帮助系统100确定对交通灯的响应(例如，灯为绿灯-移动通过枢纽，灯为红灯-在枢纽处停止)。

[0351] 在步骤3004，处理单元110(例如，经由图像对齐模块2810)可以将对应于交通灯的多个图像的区域对齐。例如，处理单元110可以使用交通灯的区域作为公共点来将多个图像对齐。该对齐可以允许处理单元110分析多个图像的公共区域。

[0352] 为了将对应于交通灯的多个图像的区域对齐，处理单元110可以选择可以允许识别关注区域的交通灯的特点或特性。在示例性实施方式中，处理单元110可以基于所确定的亮度中心(例如，交通灯的亮着的信号的中心)将对应于交通灯的多个图像的区域对齐。在一些实施方式中，亮度的中心可以通过发现例如交通灯的亮度的质心来确定。在其它实施方式中，处理单元110可以使用交通灯的颜色、形状、位置或任何其它特点或特性来将多个图像对齐。进一步，在一些实施方式中，处理单元110可以使多个图像的区域旋转，以在图像平面上将这些图像对齐(例如，匹配区域的旋转位置)。

[0353] 如上所述，经对齐的多个图像的区域可以对应于交通灯的关注区域。例如，通过基于亮度在对应于交通灯的区域处将多个图像对齐，可以识别交通灯的亮着的信号并将其放在经对齐的区域内部。亮着的信号可以向观看者指示此时可以接受的动作(例如，停止、前进、转弯等等)。因此，将对应于亮着的信号的区域对齐可以帮助处理单元110使交通灯的关注区域隔离。

[0354] 在步骤3006，处理单元110(例如，经由像素扩展模块2820)可以在经对齐的区域内部对多个图像的像素进行扩展，以获得标准尺寸(例如，相同数量的像素)。例如，处理单元110可以执行图像修改处理，在该处理中，处理单元110用多个像素来代替与多个图像的经对齐的区域相对应的每个像素。通过这种方式，一个或更多个图像的经对齐的区域可以被扩展成标准尺寸。像素的扩展可以允许像素在多个图像之间的直接比较。

[0355] 在一些实施方式中，处理单元110可以通过用像素矩阵来代替每个像素对在经对齐的区域内的多个图像的像素进行扩展。像素矩阵中的每个像素的特征是可以具有与母像素相同的值。通过这种方式，经扩展的像素还可以包括与像素处的图像的特征(例如，颜色、亮度等等)相关联的值。

[0356] 在一个实施方式中，处理单元110可以将经对齐的区域内部的每个像素扩展成3x3像素矩阵。在另一实施方式中，处理单元110可以将经对齐的区域内部的每个像素扩展成4x4像素矩阵。在其它实施方式中，可以使用其它矩阵尺寸(例如，5x5、6x6、8x8等等)。

[0357] 在一些实施方式中，像素扩展模块2820可以基于将每个像素转换成与比例因数相对应的多个像素的该比例因数(例如，2、2.5、3、4、5等等)来对像素进行扩展。像素扩展模块2820可以基于图像中对象的尺寸以及使用的标准尺寸来确定比例因数。通过这种方式，图像中的对象可以被标准化成包括相同数量的像素。

[0358] 处理单元110可以分析遍及多个图像的像素。例如，在步骤3008中，处理单元110

IPR2025-00943
Tesla EX1002 Page 988

(例如,经由图像分析模块2830)可以确定包括针对每个像素的平均像素值在内的一组平均像素值,包括经扩展的像素。如上所述,经扩展的像素包括与母像素相关联的值。像素值可以是限定所选像素并允许处理单元110确定经对齐的区域中的交通灯的特性的任何值。例如,像素值可以对应于与每个像素相关联的颜色或亮度值。在另一实施方式中,像素值可以是指示每个像素(或像素矩阵平均)是否包括阈值以上的颜色或亮度值的二进制值。处理单元110可以确定在全部的多个图像中针对每个像素的平均像素值。

[0359] 处理单元110可以使用像素值来确定经对齐的区域中交通灯的特性。例如,在步骤3010中,处理单元110可以基于一组平均像素值来确定交通灯是否包括箭头。在一个示例中,处理单元110可以将一组平均像素值与所存储的标准(例如,存储在数据库2840中的信息)进行比较,以确定像素值是否对应于箭头。在另一实施方式中,处理单元可以将平均像素值绘制地图并将其与所存储的地图进行比较,以确定像素值是否对应于箭头形状。

[0360] 在一些实施方式中,处理单元110可以使用精确度因数来确定像素值是否对应于箭头。例如,处理单元110可以确定一组平均像素值是否与所存储的标准(例如,一组像素值、所存储的图像、像素地图等等)匹配在大于阈值的精确度之内。在一些实施方式中,精确度因数可以取决于图像的质量而发生变化。例如,如果车辆200离交通灯较远和/或如果图像捕获装置122、124、126中的仅一个能够捕获交通灯的图像的话,则可以使用相对较低的精确度因数。类似地,如果车辆200离交通灯较近和/或图像捕获装置122、124、126中的超过一个(或全部)能够捕获交通灯的图像的话,则可以使用相对较高的精确度因数。在一些实施方式中,图像分析模块2830可以针对一个或更多个特征来对对象进行分类。例如,图像分析模块2830可以将对象分类为箭头形信号或圆形信号。

[0361] 通过示例性处理3000A,系统100可以被配置为识别与交通灯相关联的箭头形信号。识别箭头形信号的能力向系统100提供增强系统100的能力的其它功能。特别地,由于箭头形信号可以指示与非箭头形信号(例如,简单的圆形灯)不同的信息,系统100可以被配置为提供特定化的系统响应(例如,可以接受此时向左转弯、不可以接受此时向左转弯等等),该系统响应否则可能是不正确的(例如,绿灯指示在全部方向上前进)。在另一实施方式中,系统100可以在两个所识别的交通信号灯(例如,绿色圆形灯和红色箭头形灯)之间进行区分,以确定应当遵守哪个交通信号灯。下面将更加详细地描述用于使用箭头形信号的检测来实现该功能的示例性处理。

[0362] 图30B是示出与公开的实施方式相一致的用于基于交通灯中箭头信号的检测来引起系统响应的示例性处理3000B的流程图。在示例性实施方式中,处理单元110可以执行处理3000B,以使用处理3000A的结果来控制车辆200。例如,基于交通灯中箭头形信号的检测,处理单元110可以命令车辆200加速、转弯、停止、改变车道等等。

[0363] 在步骤3012中,处理单元110可以识别交通信号灯。例如,处理单元110可以在由一个或更多个图像捕获装置122、124、126捕获的一个或更多个图像中识别交通信号灯。在一些情况下,可以在图像中识别超过一个的交通信号灯(例如,绿色圆形灯和红色箭头形灯)。

[0364] 在步骤3014,处理单元110可以确定所识别的交通信号灯的颜色。处理单元110可以使用颜色识别处理来分析交通灯的多个图像中的一个或更多个图像,以确定交通信号灯的颜色并对其进行分类。例如,处理单元110可以将图像数据与所存储的标准进行匹配,以从多种颜色可能性(例如,绿色、红色、黄色、橙色等等)中确定颜色。

IPR2025-00943
Tesla EX1002 Page 989

[0365] 在一个实施方式中，处理单元110可以识别包括指示像素是交通信号灯的部分形状的像素值在内的像素的颜色。在另一实施方式中，一个或更多个图像捕获装置122、124、126可以使用清晰像素和红色(或其它颜色)像素来帮助进行颜色区分。例如，在一些实施方式中，当图像中存在红灯时，可以使用红色像素和清晰像素，而当存在绿灯(或其它颜色的)灯时，仅可以使用清晰像素而忽略红色像素(例如，定义为黑色、无颜色、不是红色等等)。例如，当灯是红灯时，红色像素可以与颜色匹配并且被“用来”确定该颜色。作为另一示例，当灯是绿灯时，可以使用清晰像素来识别颜色和/或忽略掉红色像素以将颜色限定为不是红色。通过这种方式，处理单元110可以取决于被用来限定像素内部的灯的值的像素的类型来在图像中的颜色之间进行区分。进一步，由于图像可以与其它图像进行对齐并进行平均化以实现更高分辨率的图像，可以基于来自其它图像的信息来提供所忽略红色像素(或缺失值的其它像素)的值。

[0366] 应当理解，上述内容是示例并且可以使用其它颜色识别处理和/或技术。在一些实施方式中，处理单元110可以被配置为确定箭头和/或交通信号的一个或更多个其它特性。例如，处理单元110可以确定箭头是闪烁的还是稳定的。在另一示例中，处理单元110可以确定箭头是否最近从其它信号变来。例如，多个图像可以包括两个所识别的信号，之前信号和当前信号(例如，所识别的箭头)。这些其它特性可以为处理单元110提供可以有助于确定如何解释箭头形信号的附加信息。

[0367] 在步骤3018中，处理单元110可以基于箭头的颜色(和/或箭头和/或交通信号的其它特性)来产生一个或更多个系统响应。在一个示例中，系统响应可以对应于箭头的指示(例如，绿色箭头-可以接受转弯、红色箭头-不可以接受转弯等等)。在另一示例中，系统响应可以对应于特定交通信号由于其是箭头而不适用的指示(例如，车辆正在直行行驶)。系统响应可以是被配置为改变或继续车辆200的操作参数的导航响应。例如，处理单元110可以向车辆的驾驶员提供通知(例如，经由用户接口170的可听声音或可视指示符)，施加车辆制动，中止巡航控制，和/或发起自动转弯机动。例如，处理单元110可以向节流系统220、制动系统230和转向系统240中的一个或更多个提供控制信号以对车辆200进行导航。在一些实施方式中，处理单元110可以继续对图像进行分析，以确定是否应当中断系统响应(例如，箭头不再是绿色-中止巡航控制并施加制动)。

[0368] 通过示例性公开的处理3000A和3000B，系统100可以具有提供车辆200的自主控制的增强的功能。特别地，系统100可以被配置为基于由交通信号提供的详细信息来对交通信号做出反应，并且提供允许车辆200的安全操作的恰当的系统响应。进一步，尽管已经在箭头形信号的语境下描述了上述处理，应当理解的是，其一个或更多个步骤可以被用于相似处理中，以确定对其它对象和/或信号(诸如行人信号、红灯不转弯信号、固定道路标志、街道标志、交通咨询牌、广告牌等等)的系统响应。

[0369] 出于图示的目的，前面的描述已呈现。并非穷尽且并不限于公开的具体形式或实施方式。考虑到说明书和公开的实施方式的实践，修改和改写对本领域技术人员将是明显的。另外，虽然公开的实施方式的各方面描述为存储在存储器中，但是本领域技术人员将了解到，这些方面也可以存储在其它类型的诸如二级存储装置的计算机可读介质上，例如，硬盘或CD ROM，或其它形式的RAM或ROM、USB介质、DVD、蓝光或其它光驱介质。

[0370] 基于书面描述和公开方法的计算机程序在有经验的开发人员的技术内。各种程序

IPR2025-00943
Tesla EX1002 Page 990

或程序模块可以使用本领域技术人员公知的任何技术来创建或者可以结合现有软件设计。例如，程序段或程序模块可以设计或借助.Net Framework、.Net Compact Framework(及相关语言，诸如Visual Basic、C，等等)、Java、C++、Objective-C、HTML、HTML/AJAX组合、XML或包括Java小程序的HTML。

[0371] 而且，虽然说明性实施方式已在本文中描述，具有等同元件、修改、省略、组合(例如，跨各实施方式的方面)、改写和/或变更的任何及所有实施方式的范围将基于本公开被本领域技术人员了解。权利要求书中的限制应解译宽泛地基于权利要求书中采用的语言，并不限于本说明书或本申请审查期间描述的示例。示例应被解释为非排他性的。此外，公开方法中的步骤能以任何方式修改，包括重新排序各步骤和/或插入或删除步骤。因此，预期说明书和示例考虑为仅是说明性的，真正的范围和精神由随附权利要求书及其完整的等同范围指示。

IPR2025-00943
Tesla EX1002 Page 991

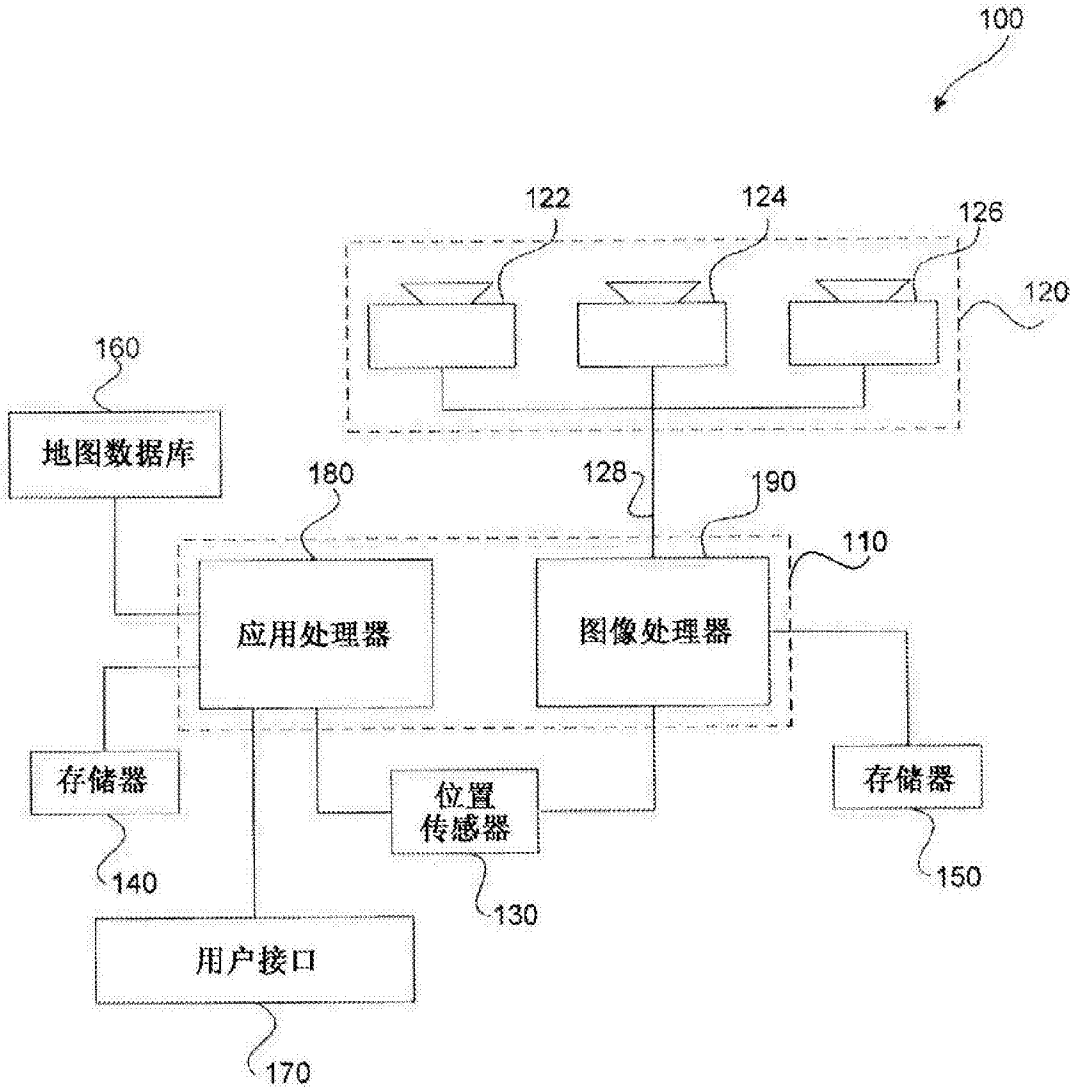


图1

IPR2025-00943
Tesla EX1002 Page 992

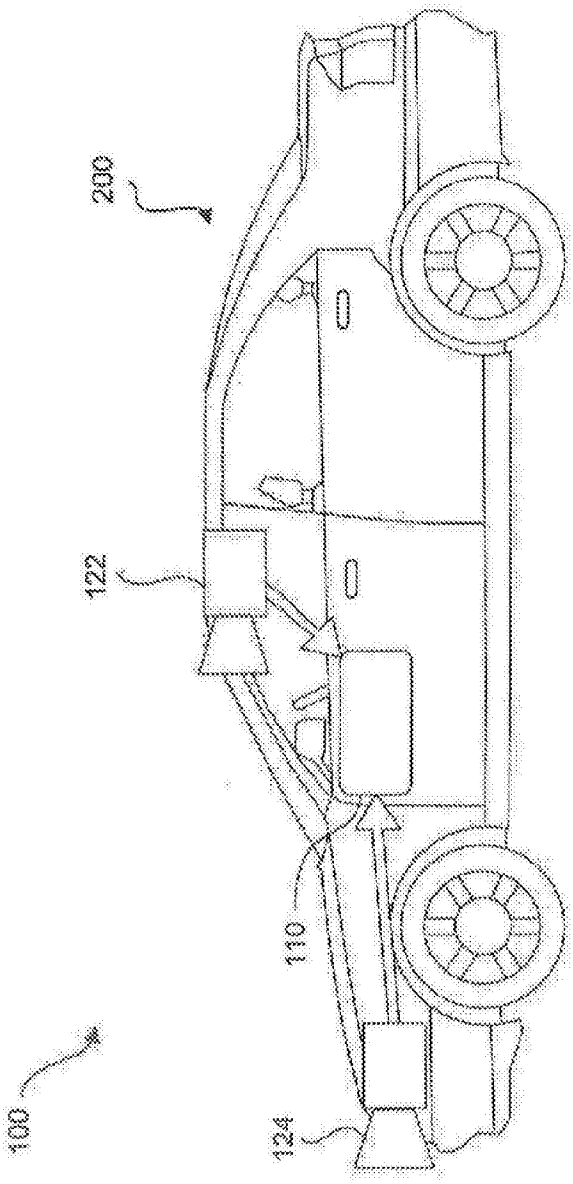


图2A

IPR2025-00943
Tesla EX1002 Page 993

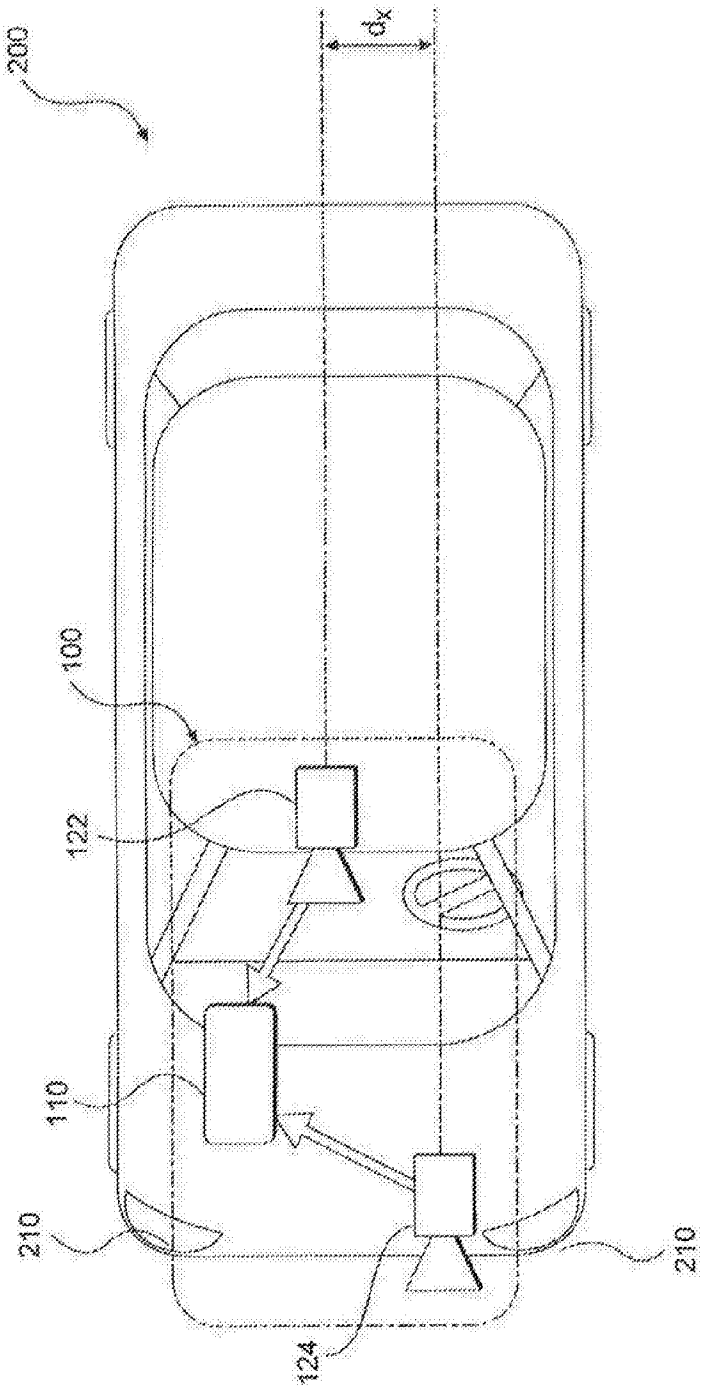


图2B

IPR2025-00943
Tesla EX1002 Page 994

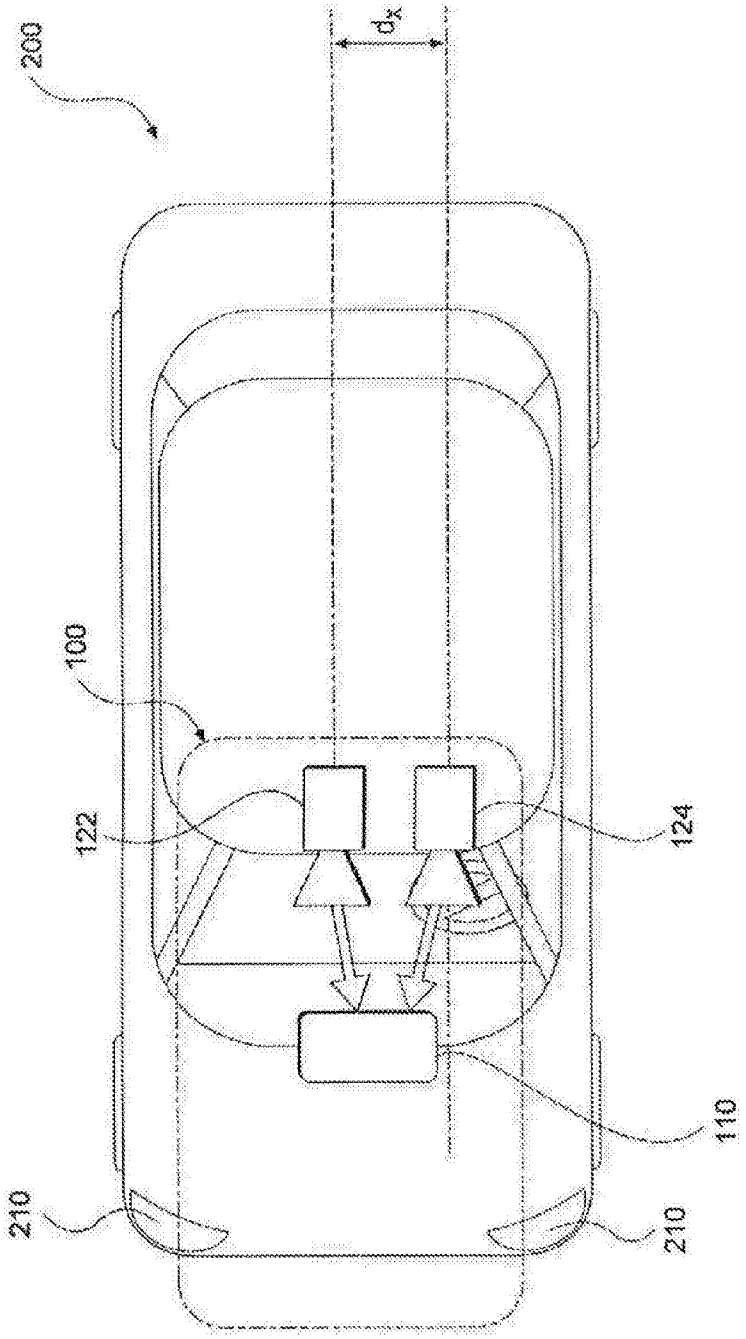


图2C

IPR2025-00943
Tesla EX1002 Page 995

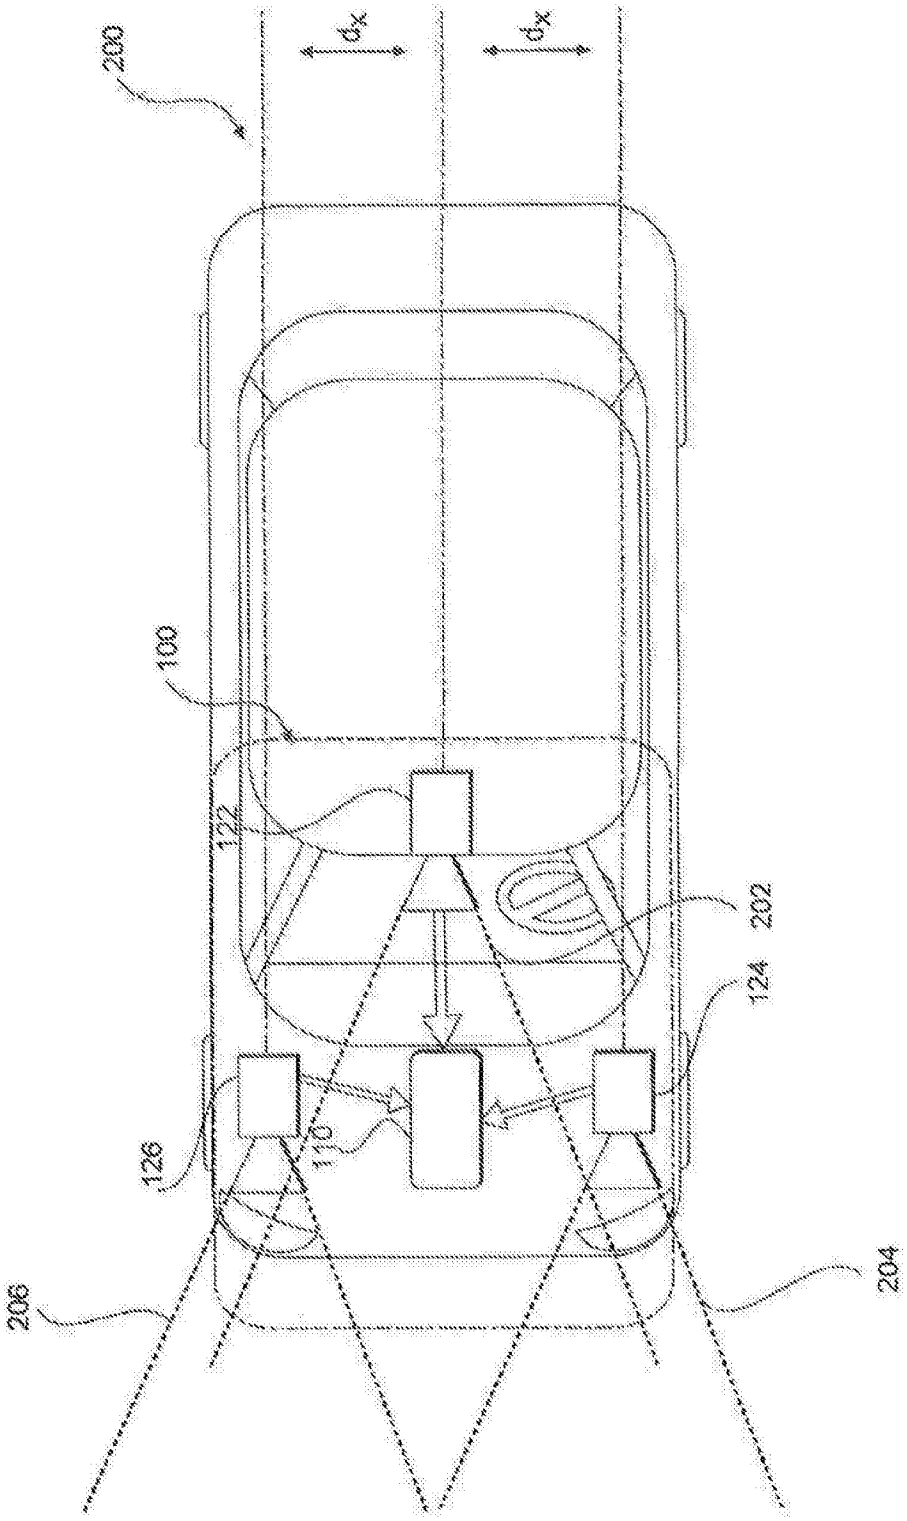


图2D

IPR2025-00943
Tesla EX1002 Page 996

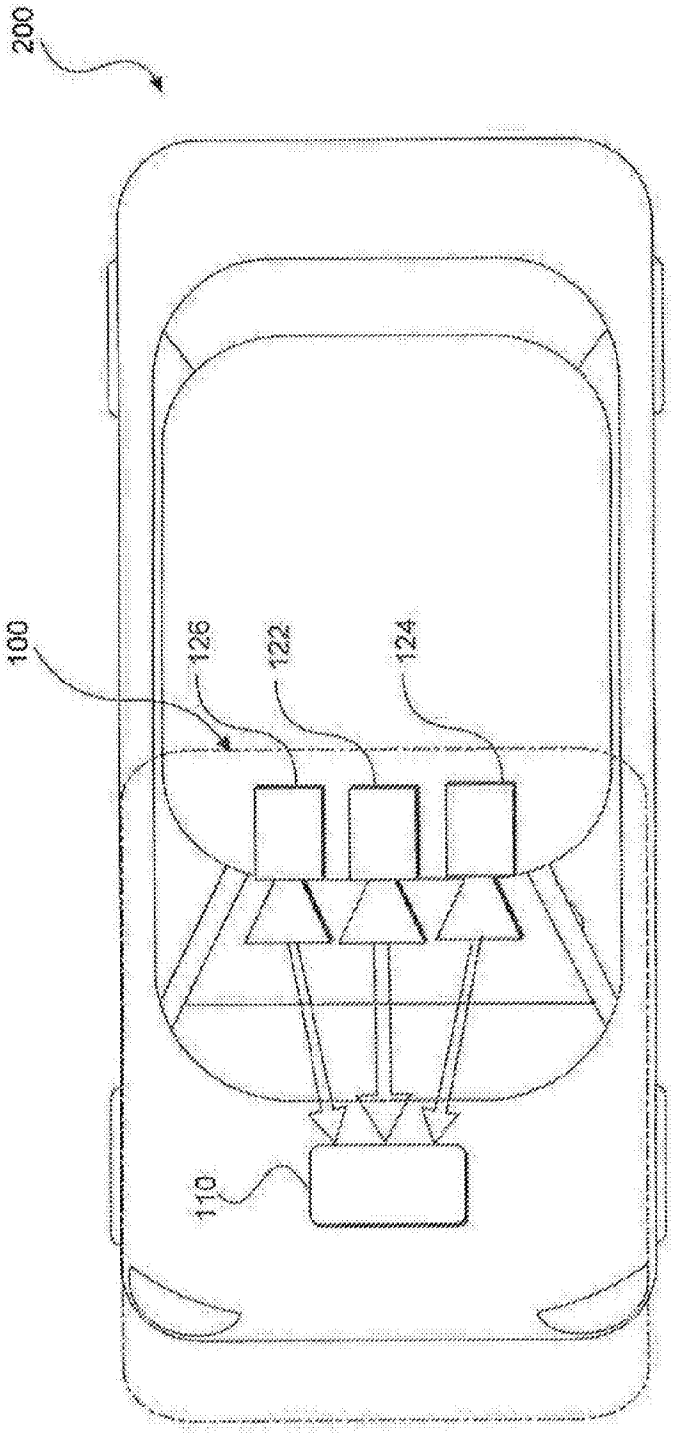


图2E

IPR2025-00943
Tesla EX1002 Page 997

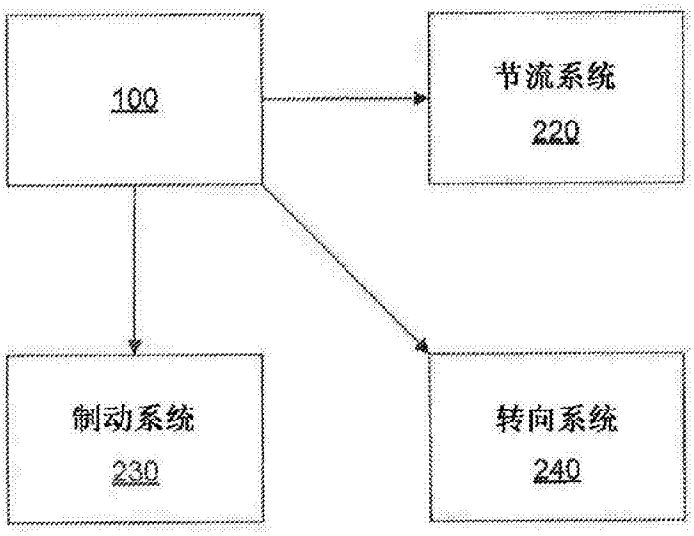


图2F

IPR2025-00943
Tesla EX1002 Page 998

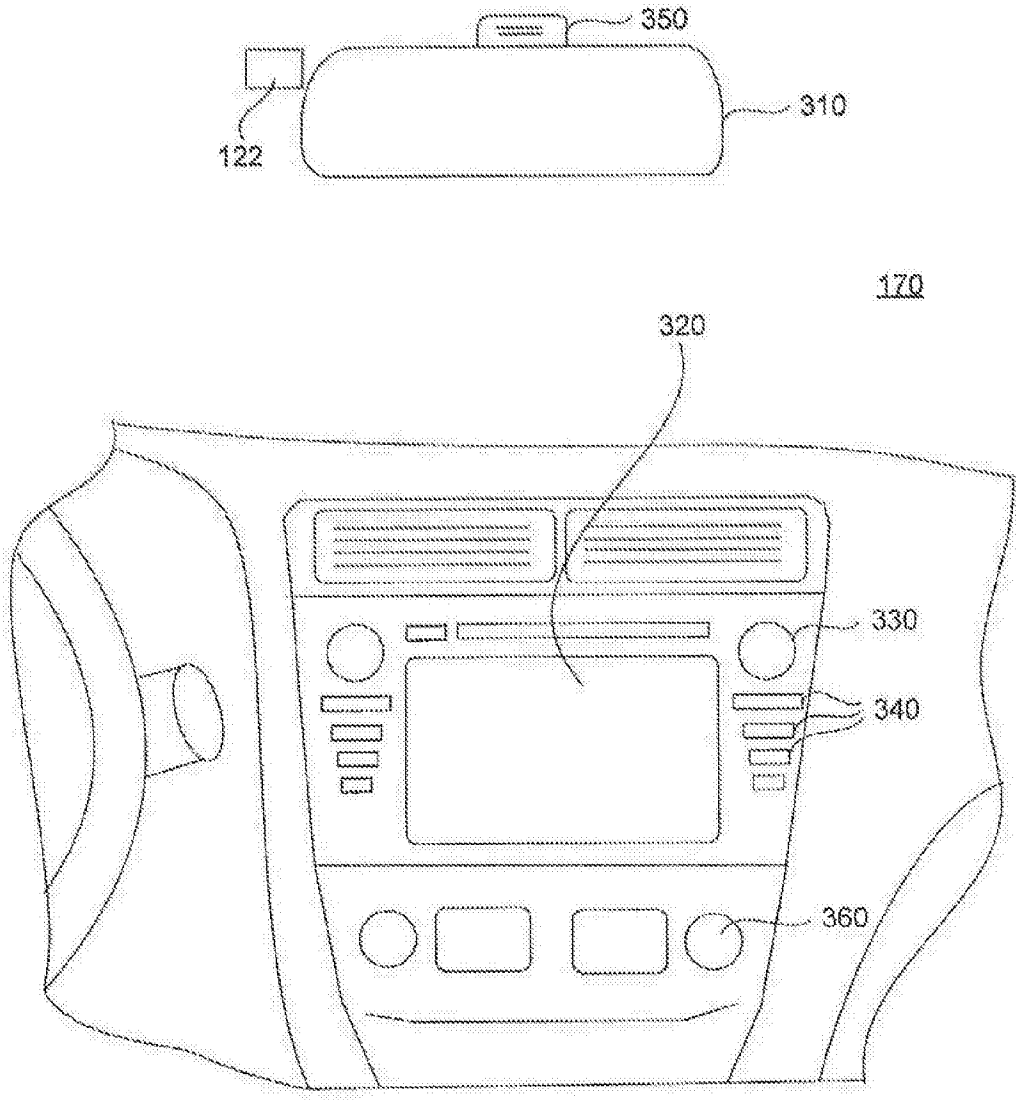


图3A

IPR2025-00943
Tesla EX1002 Page 999

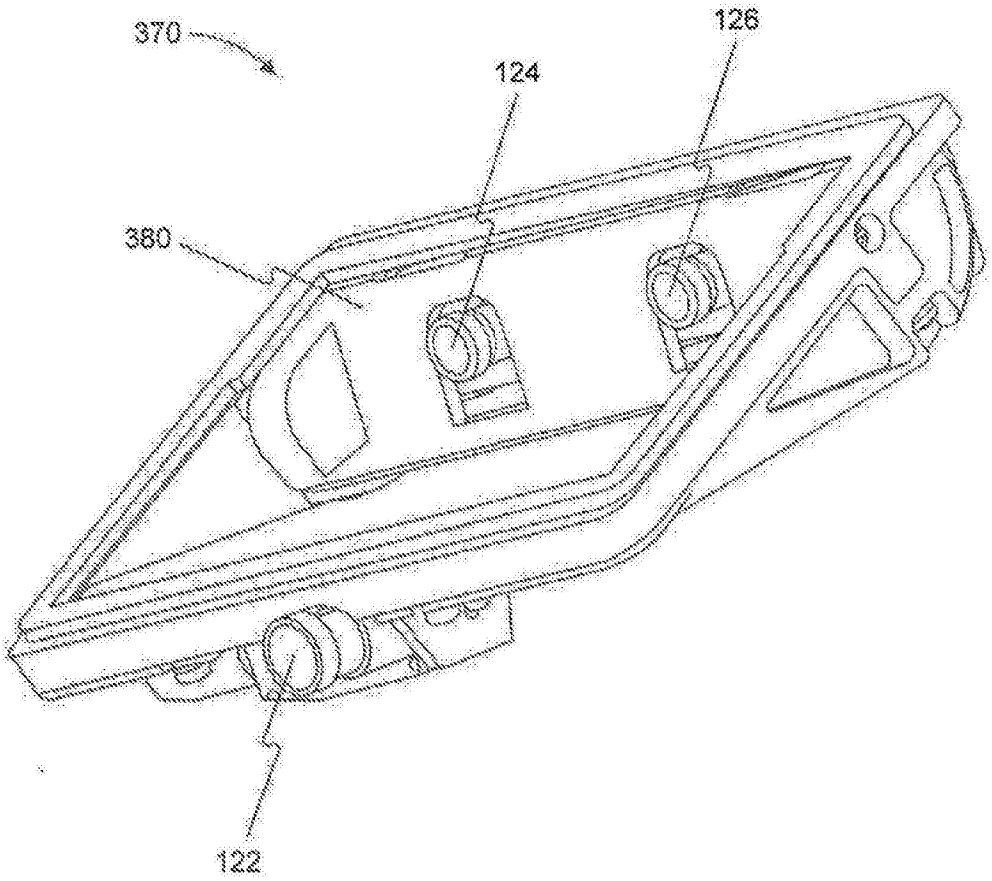


图3B

IPR2025-00943
Tesla EX1002 Page 1000

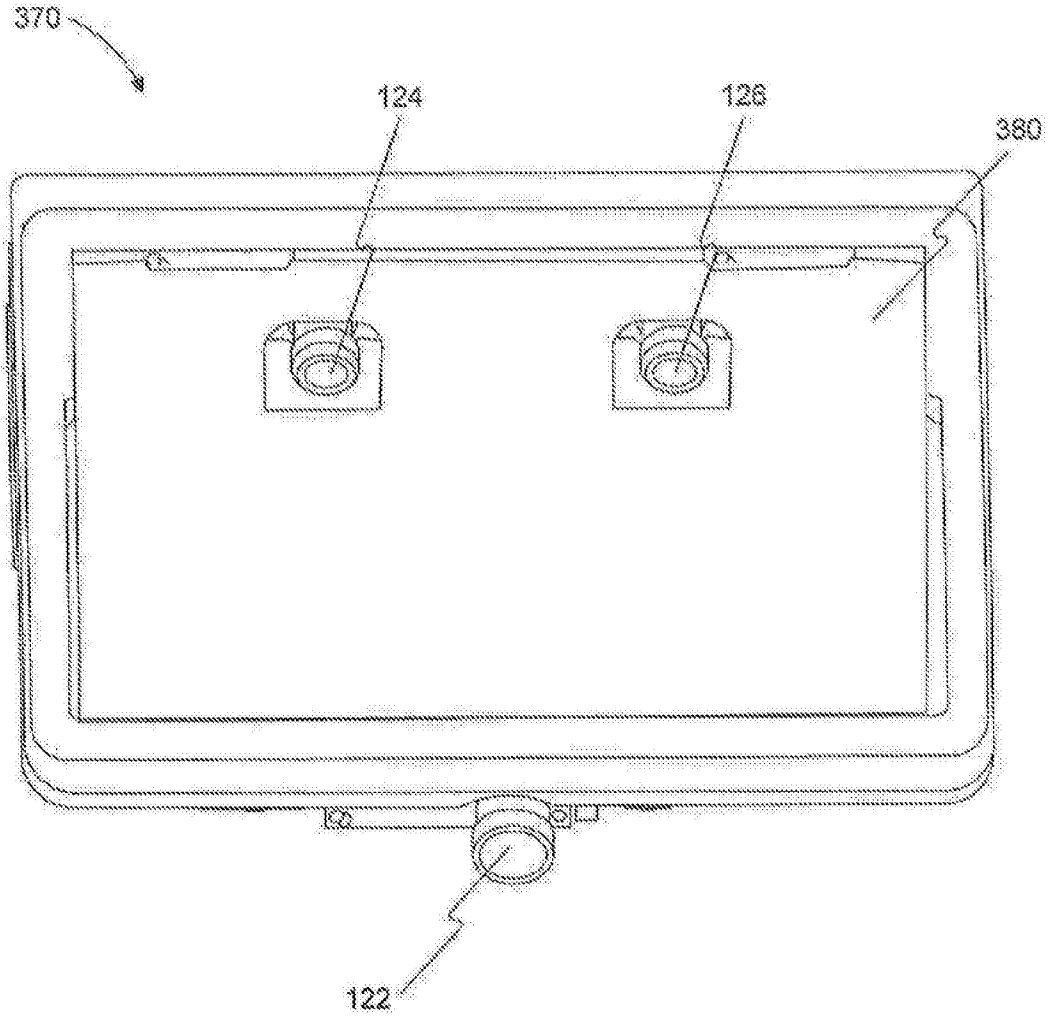


图3C

IPR2025-00943
Tesla EX1002 Page 1001

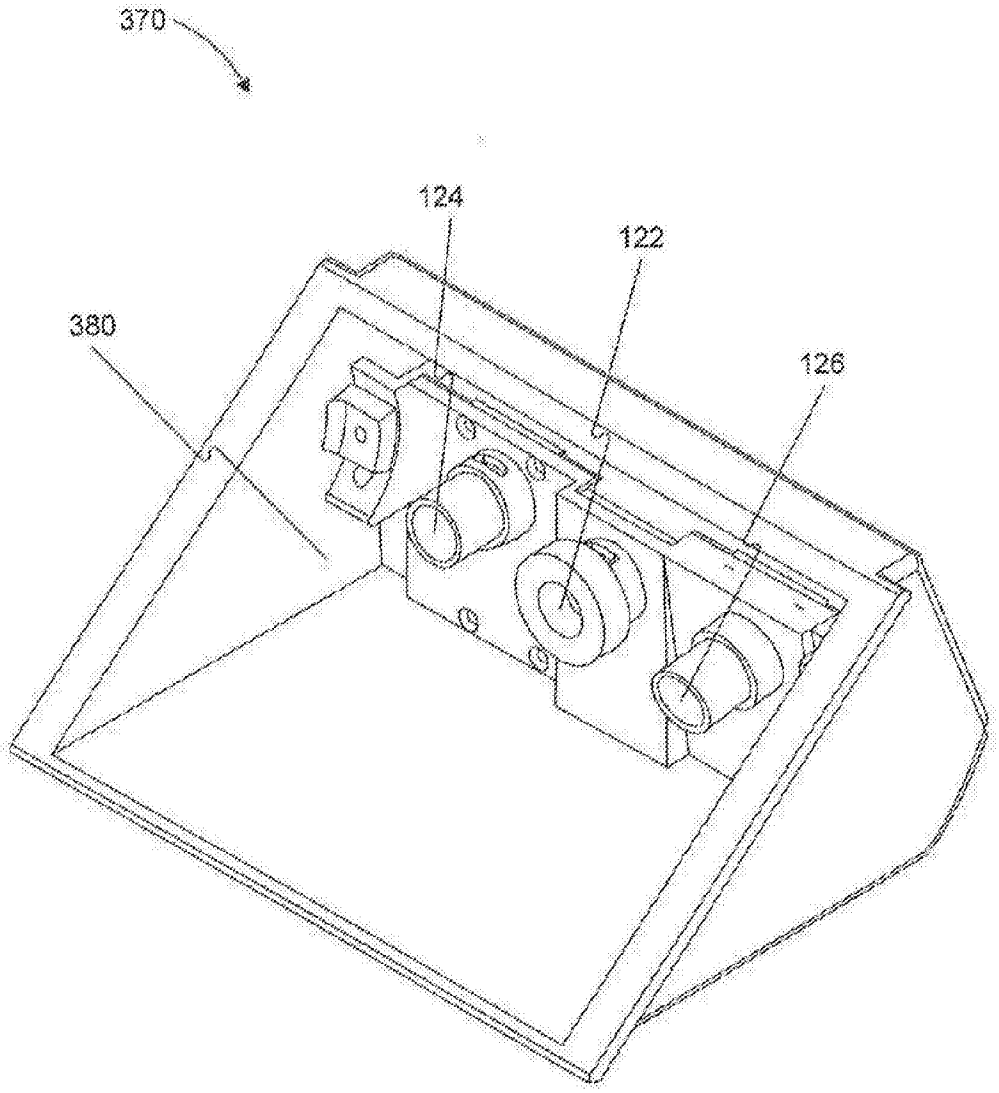


图3D

IPR2025-00943
Tesla EX1002 Page 1002

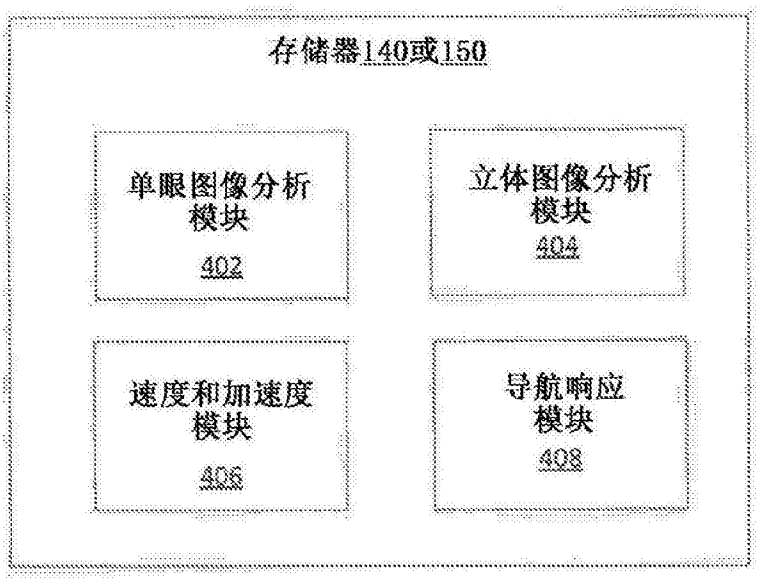


图4

IPR2025-00943
Tesla EX1002 Page 1003

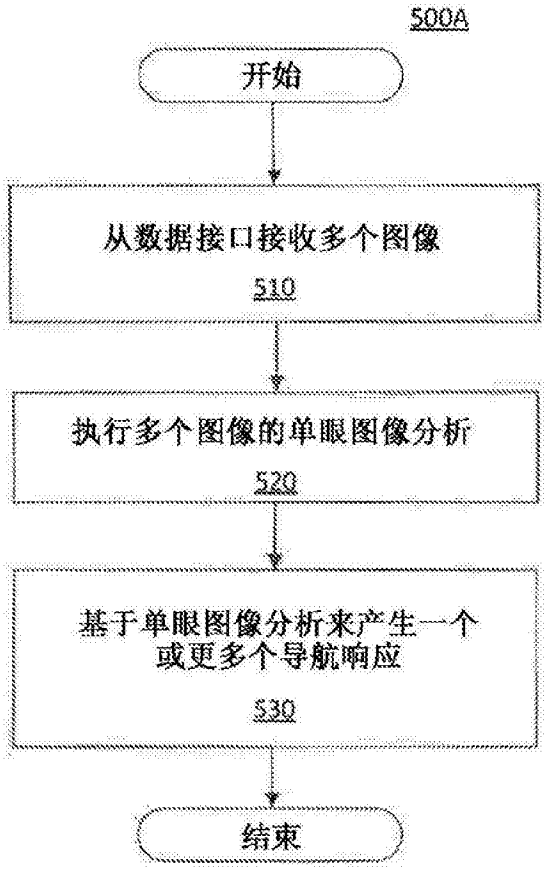


图5A

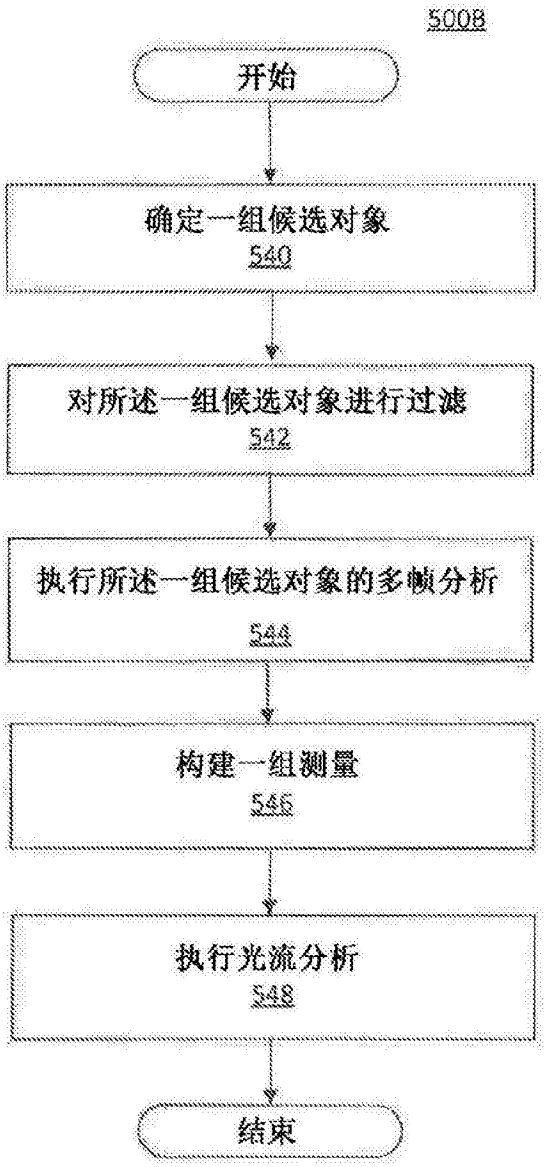


图5B

IPR2025-00943
Tesla EX1002 Page 1004

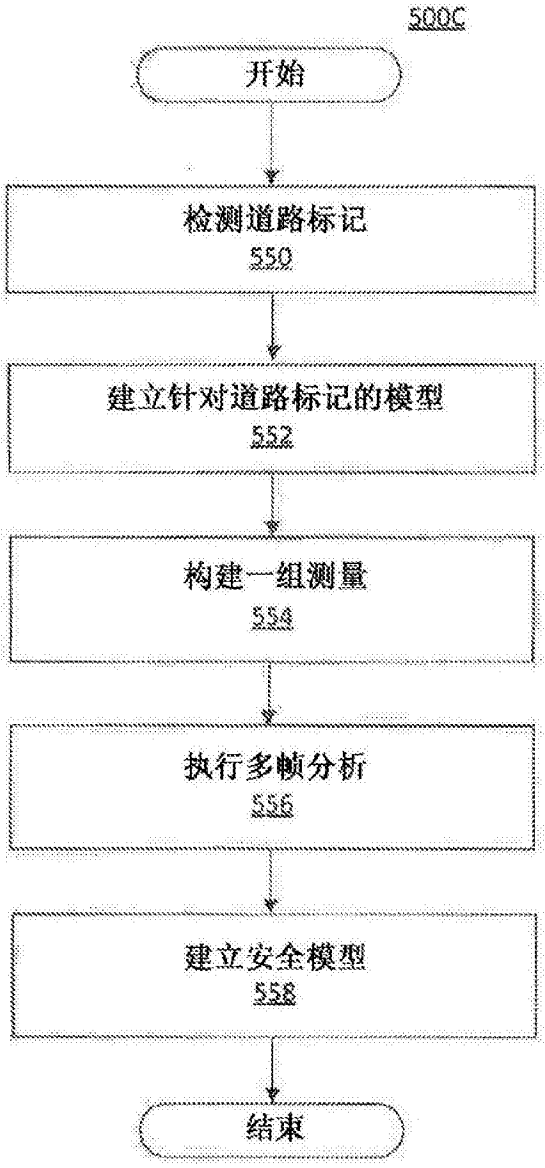


图5C

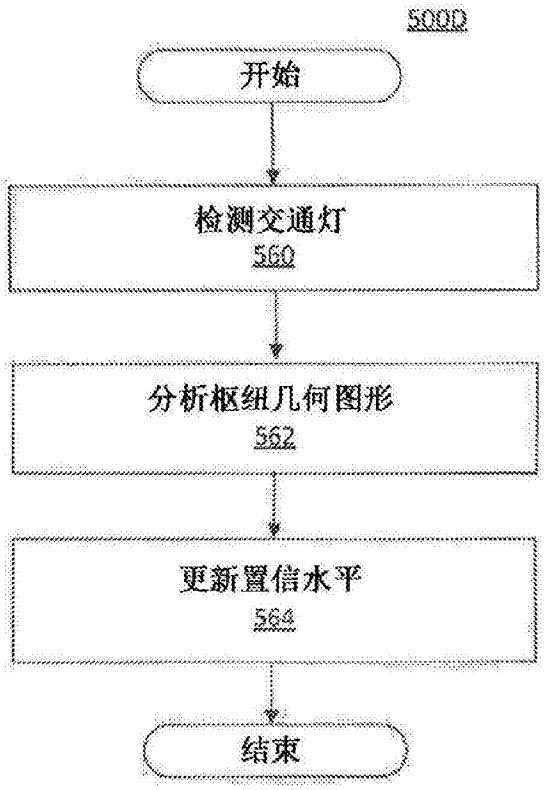


图5D

IPR2025-00943
Tesla EX1002 Page 1005

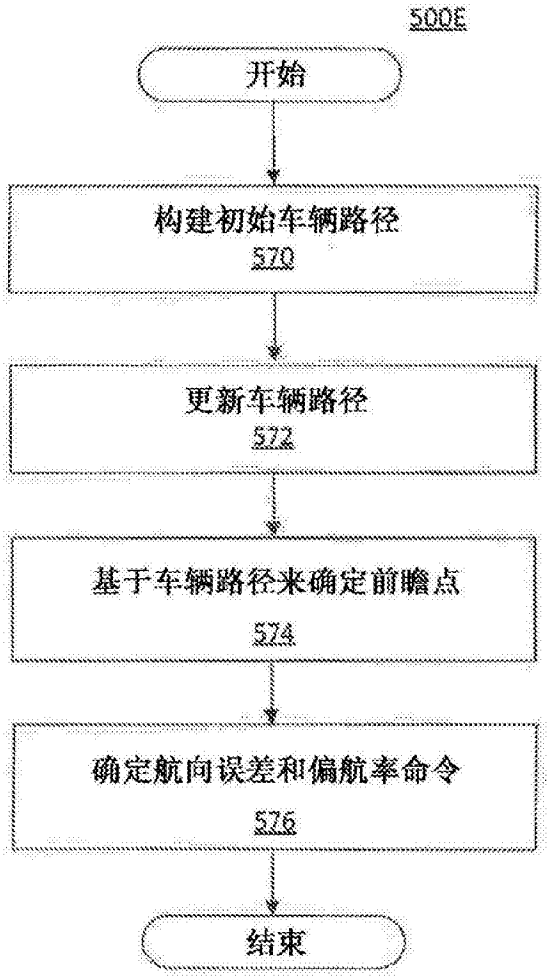


图5E

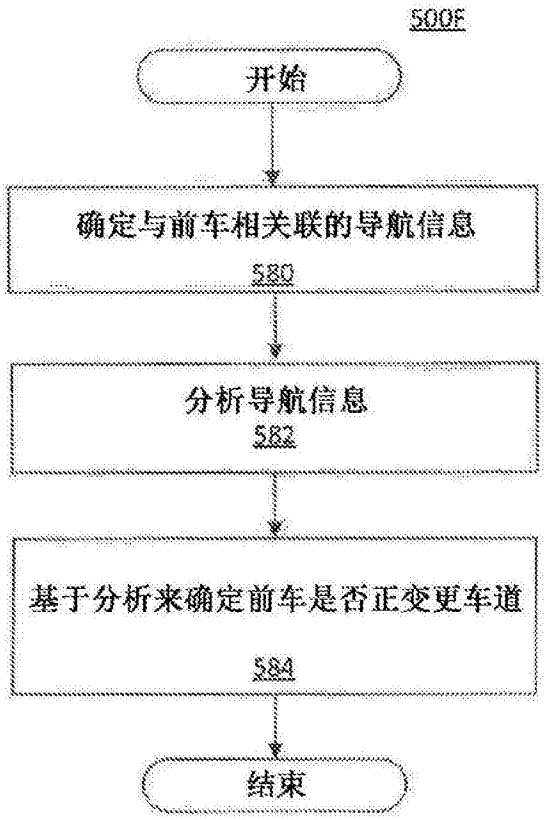


图5F

IPR2025-00943
Tesla EX1002 Page 1006

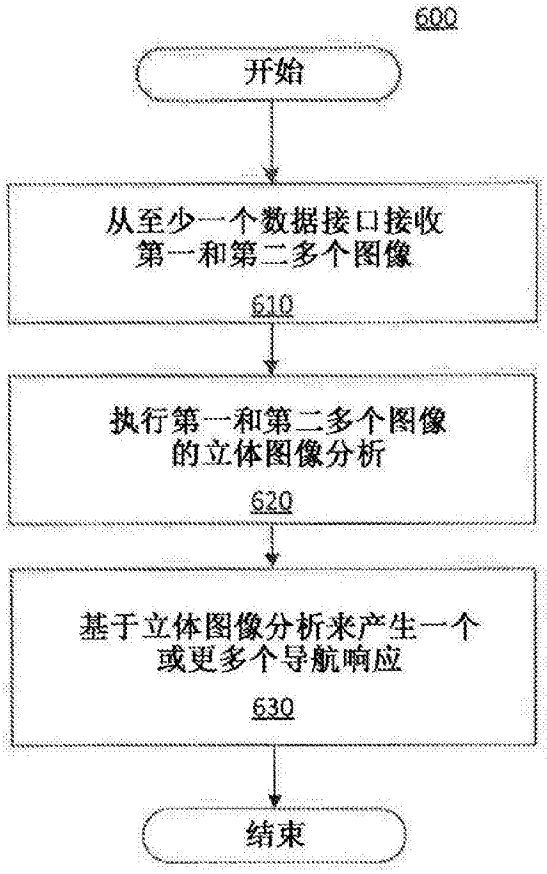


图6

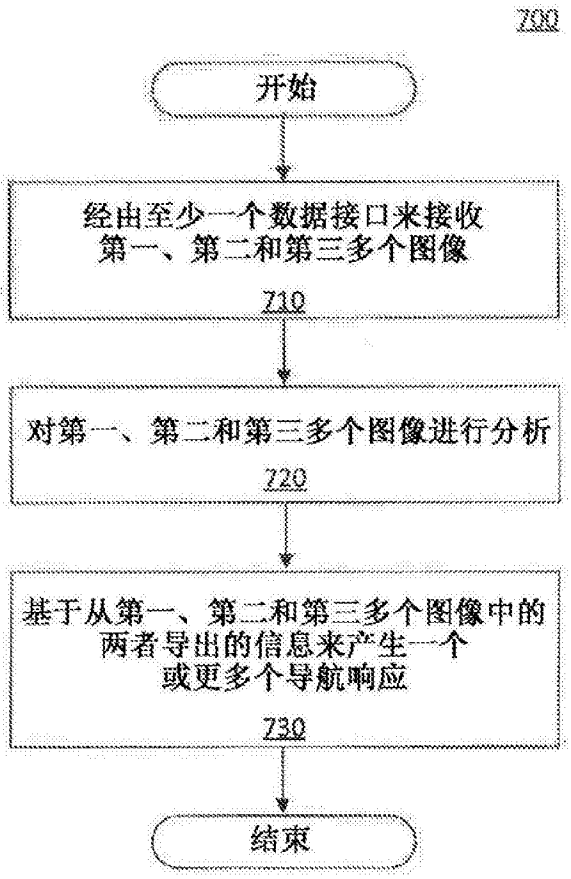


图7

IPR2025-00943
Tesla EX1002 Page 1007

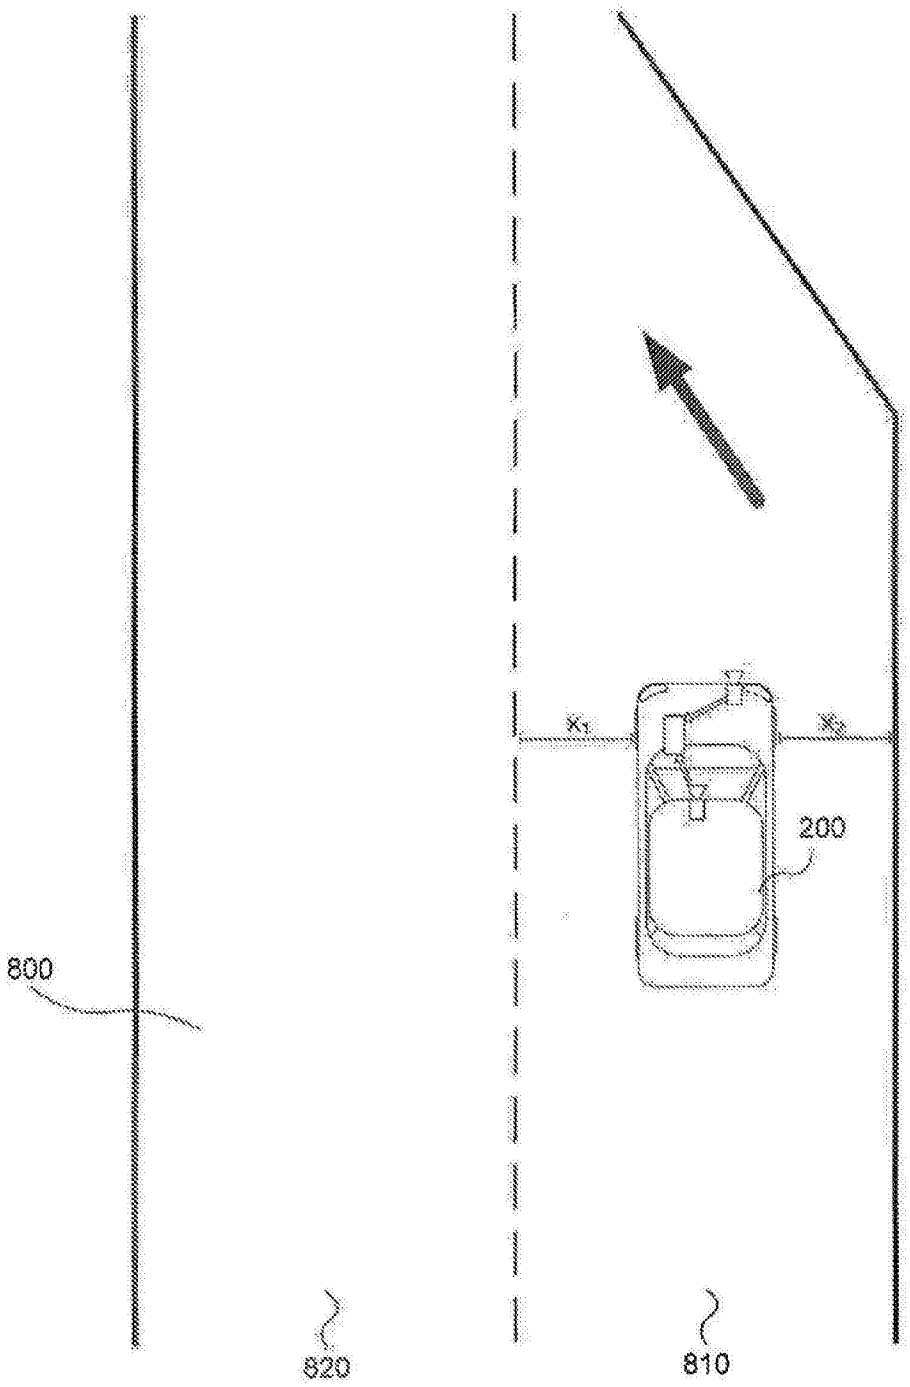


图8

IPR2025-00943
Tesla EX1002 Page 1008

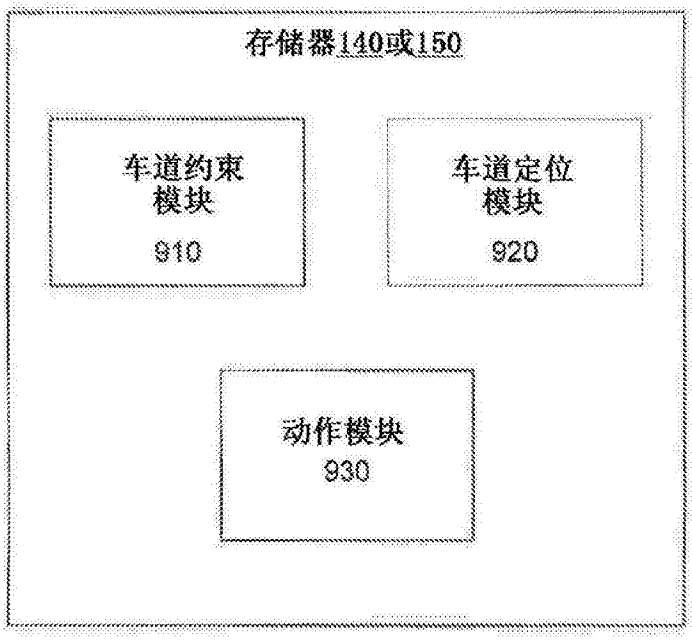


图9

IPR2025-00943
Tesla EX1002 Page 1009

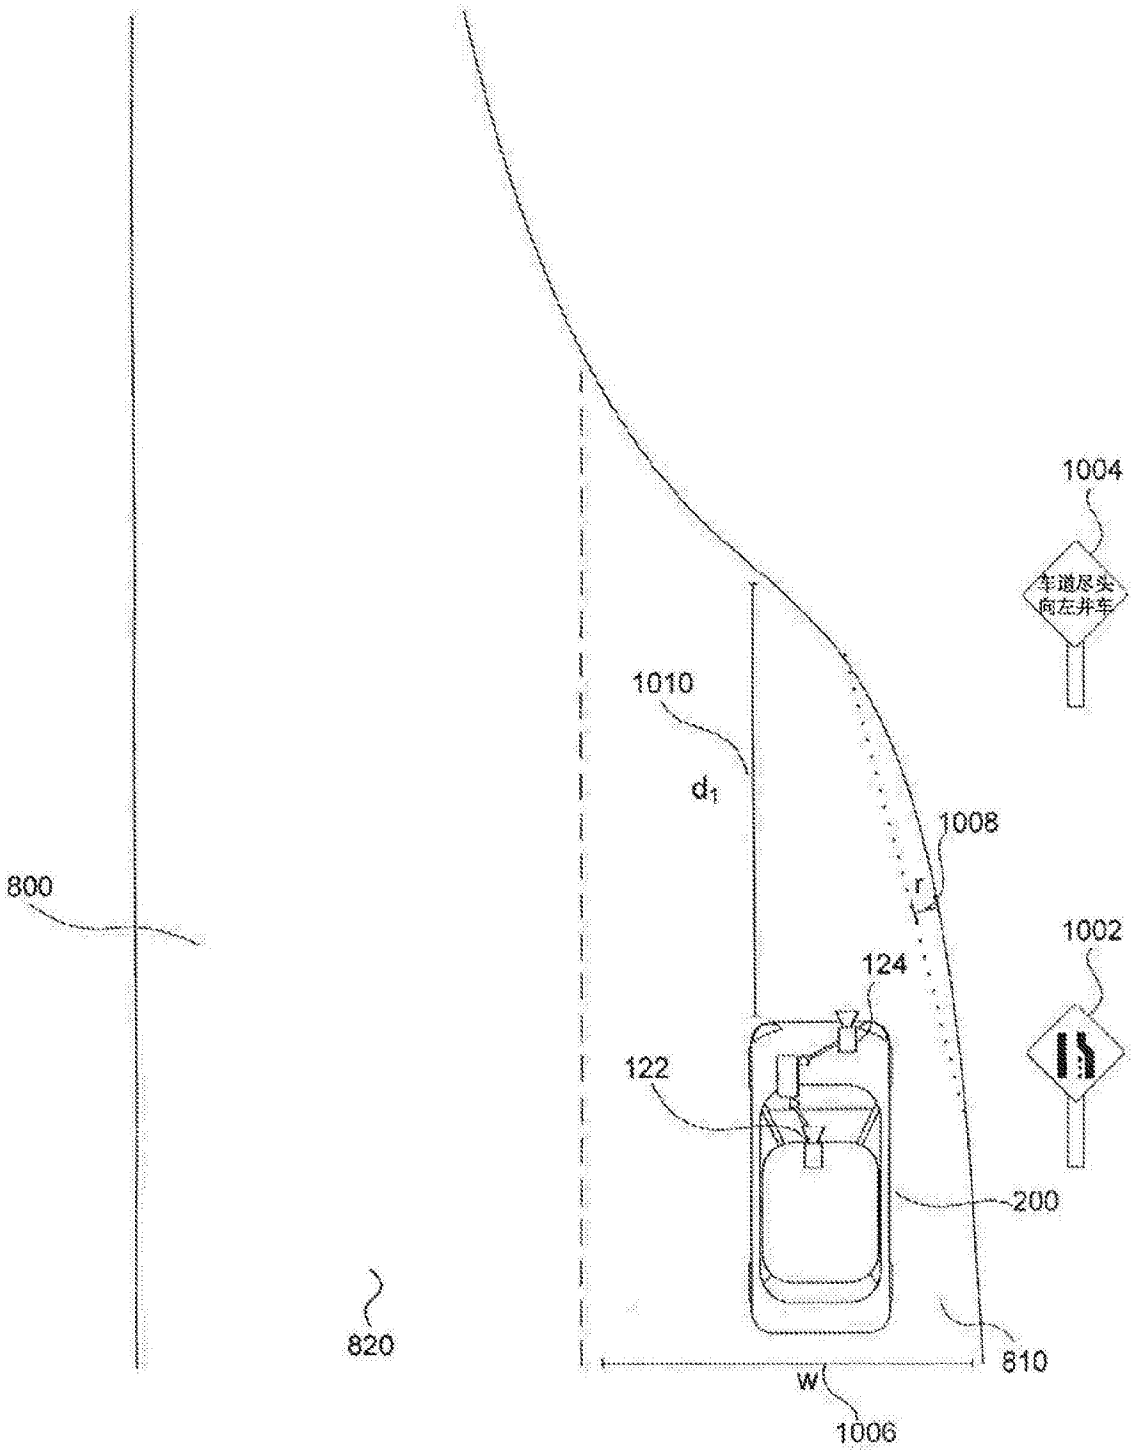


图10A

IPR2025-00943
Tesla EX1002 Page 1010

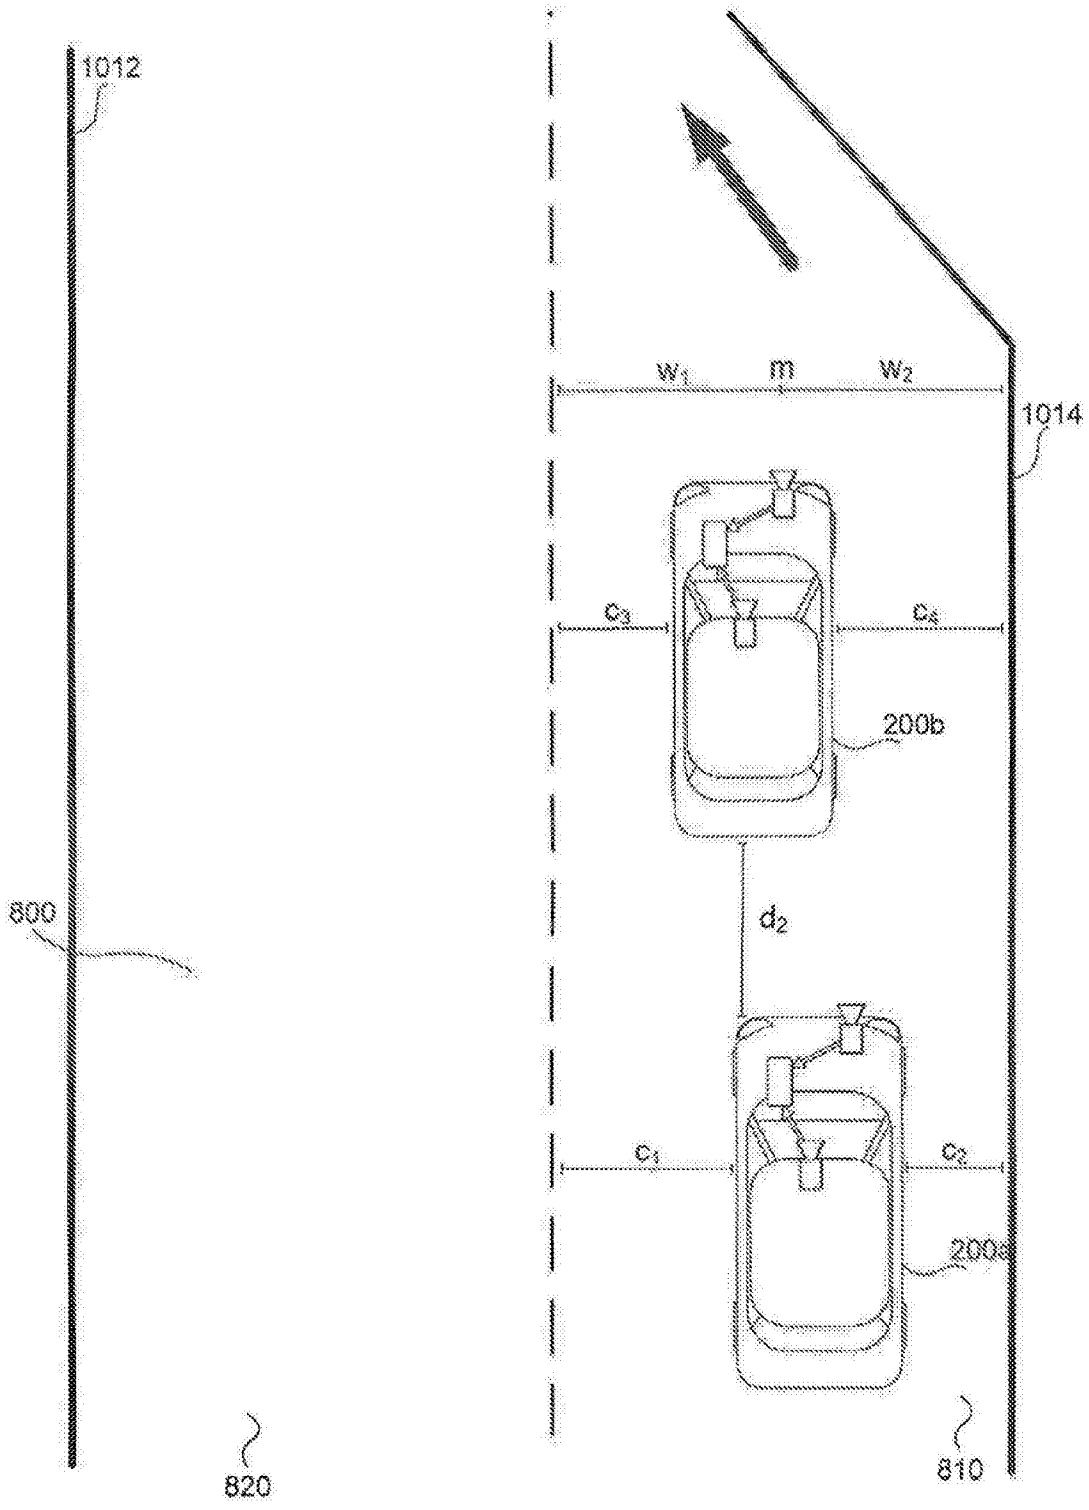


图10B

IPR2025-00943
Tesla EX1002 Page 1011

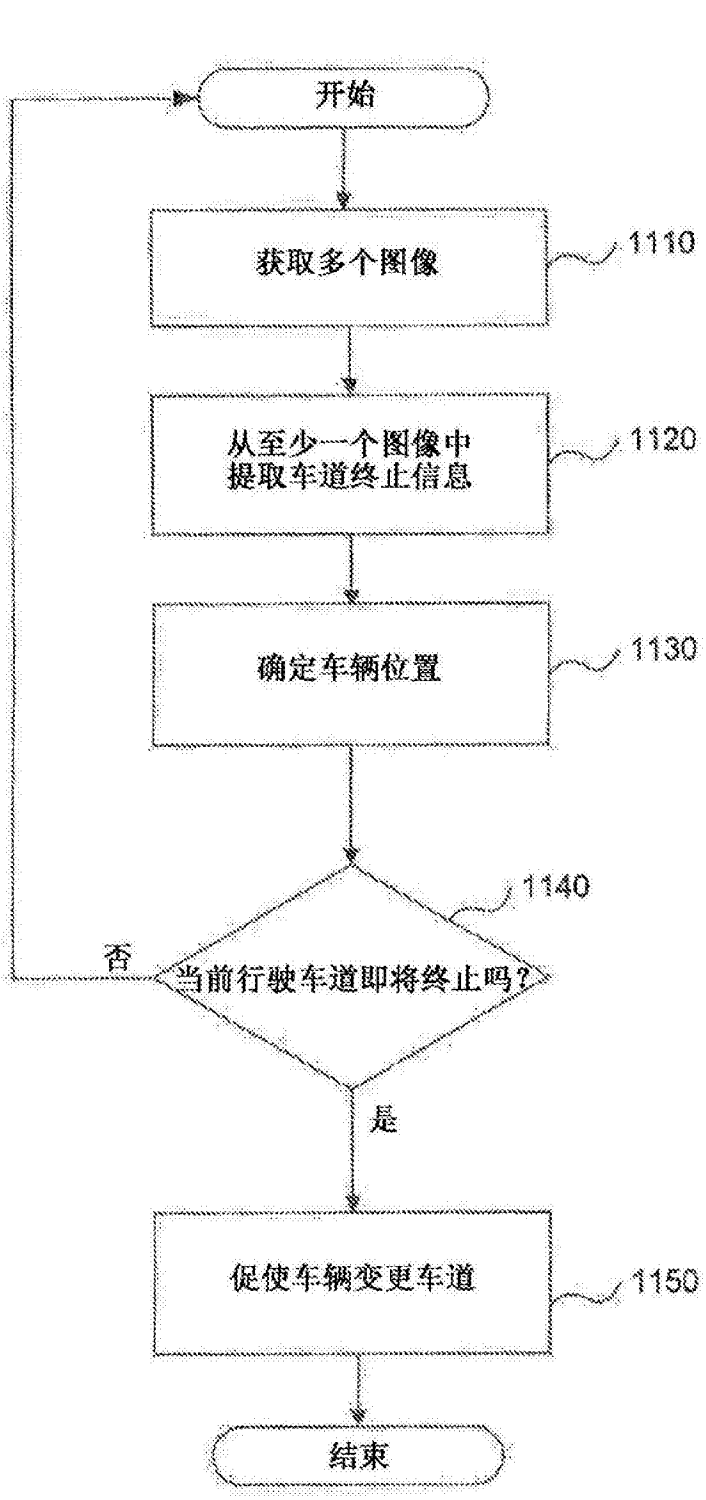


图11

IPR2025-00943
Tesla EX1002 Page 1012

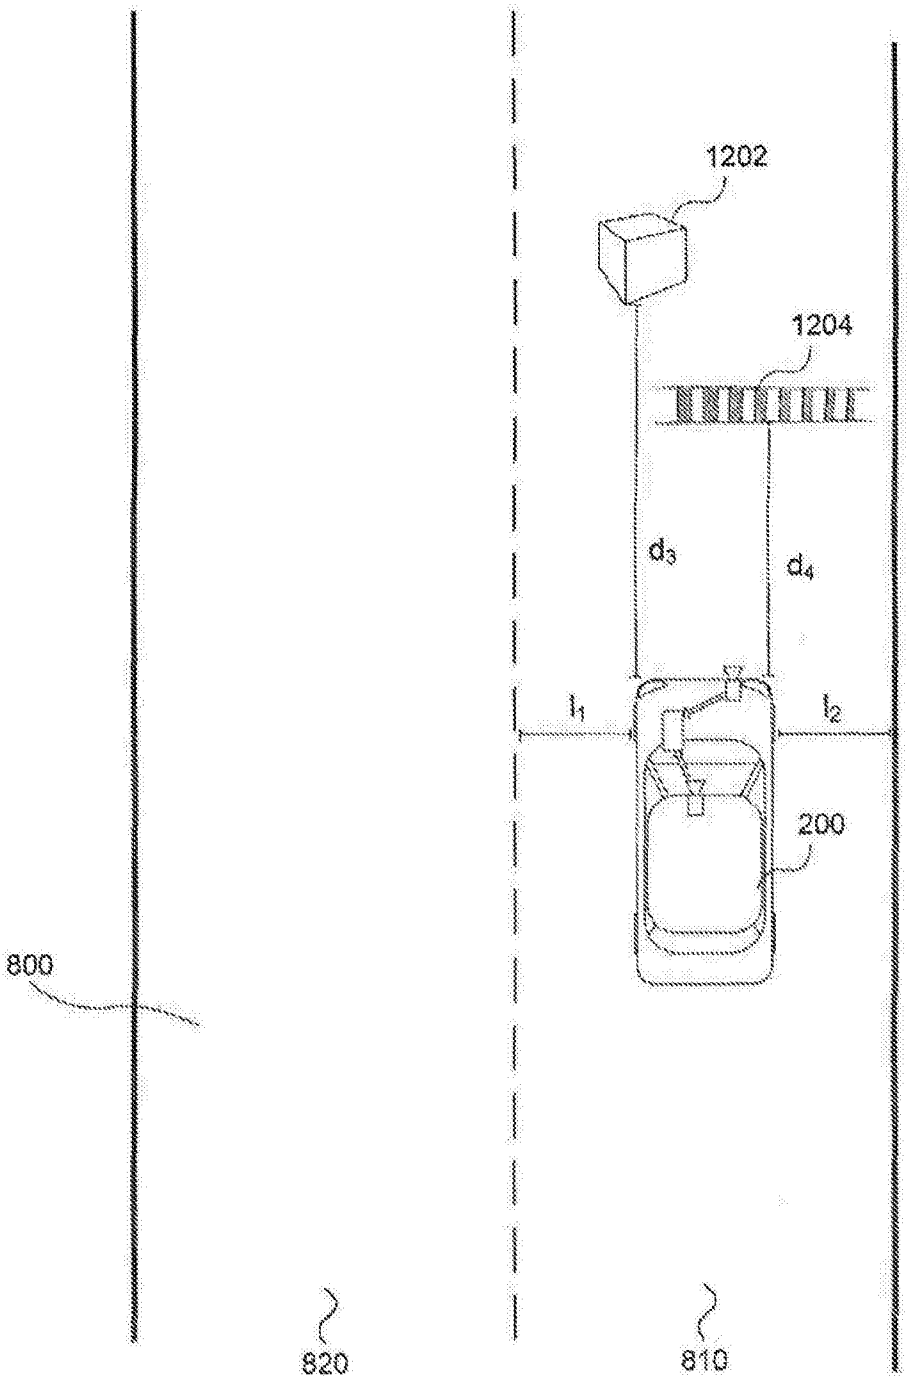


图12

IPR2025-00943
Tesla EX1002 Page 1013

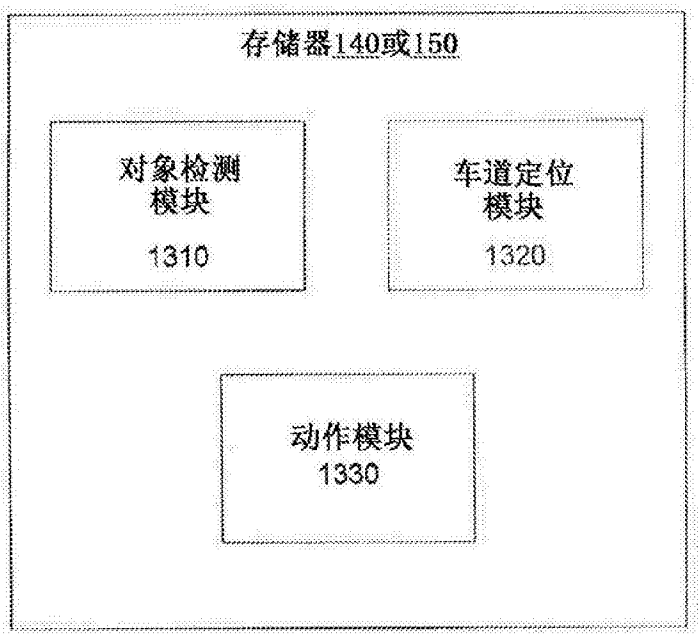


图13

IPR2025-00943
Tesla EX1002 Page 1014

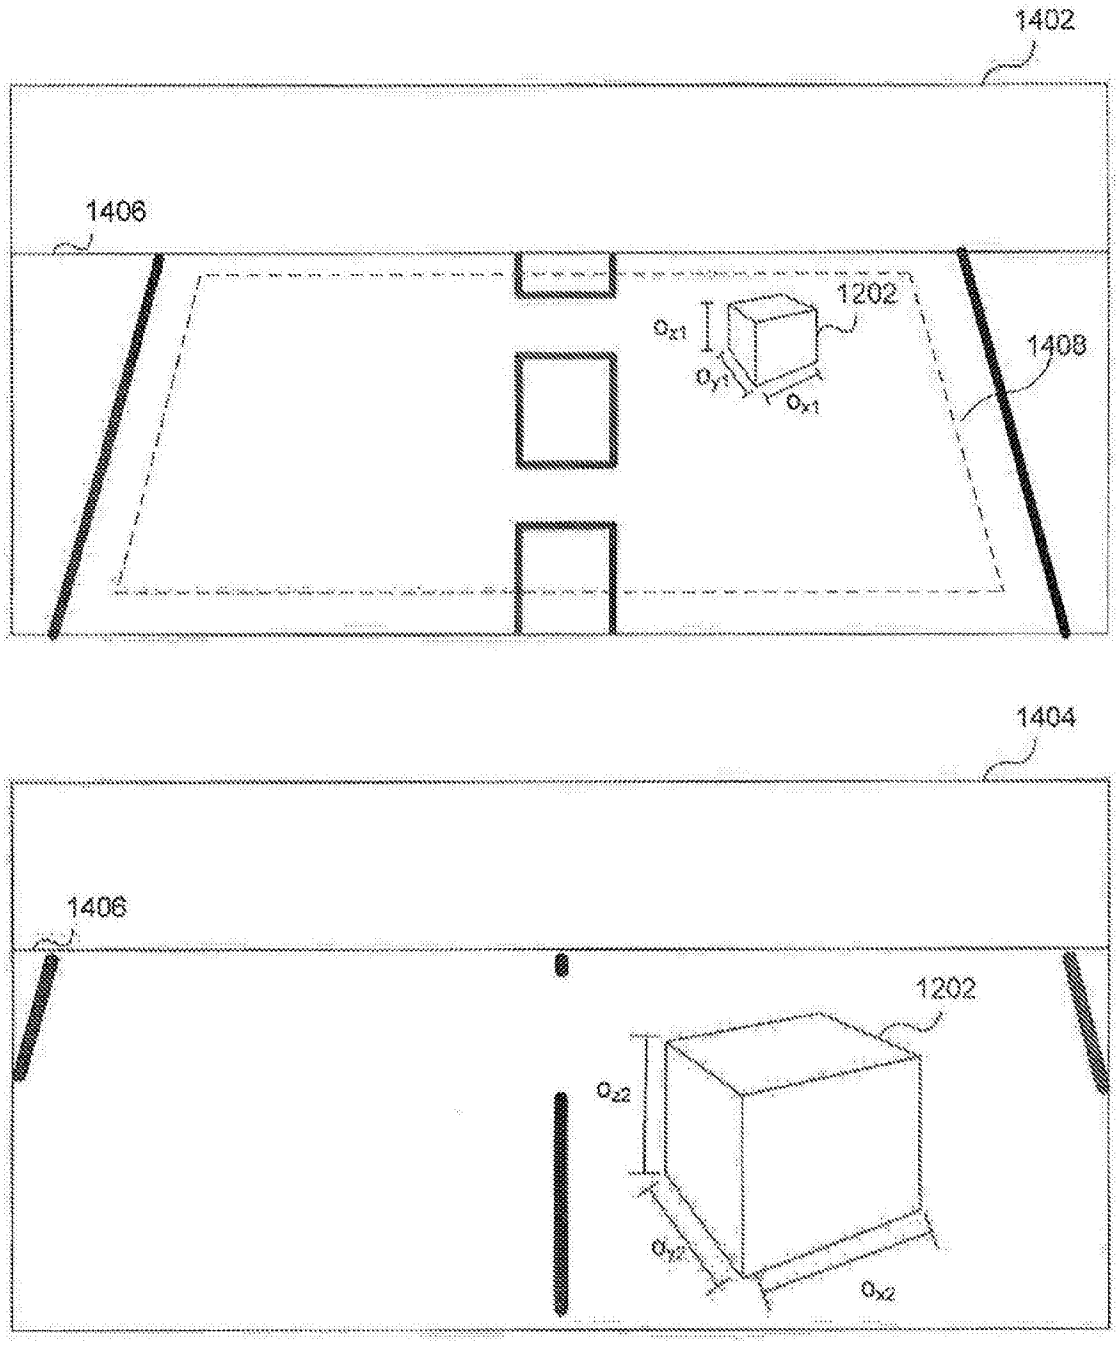


图14A

IPR2025-00943
Tesla EX1002 Page 1015

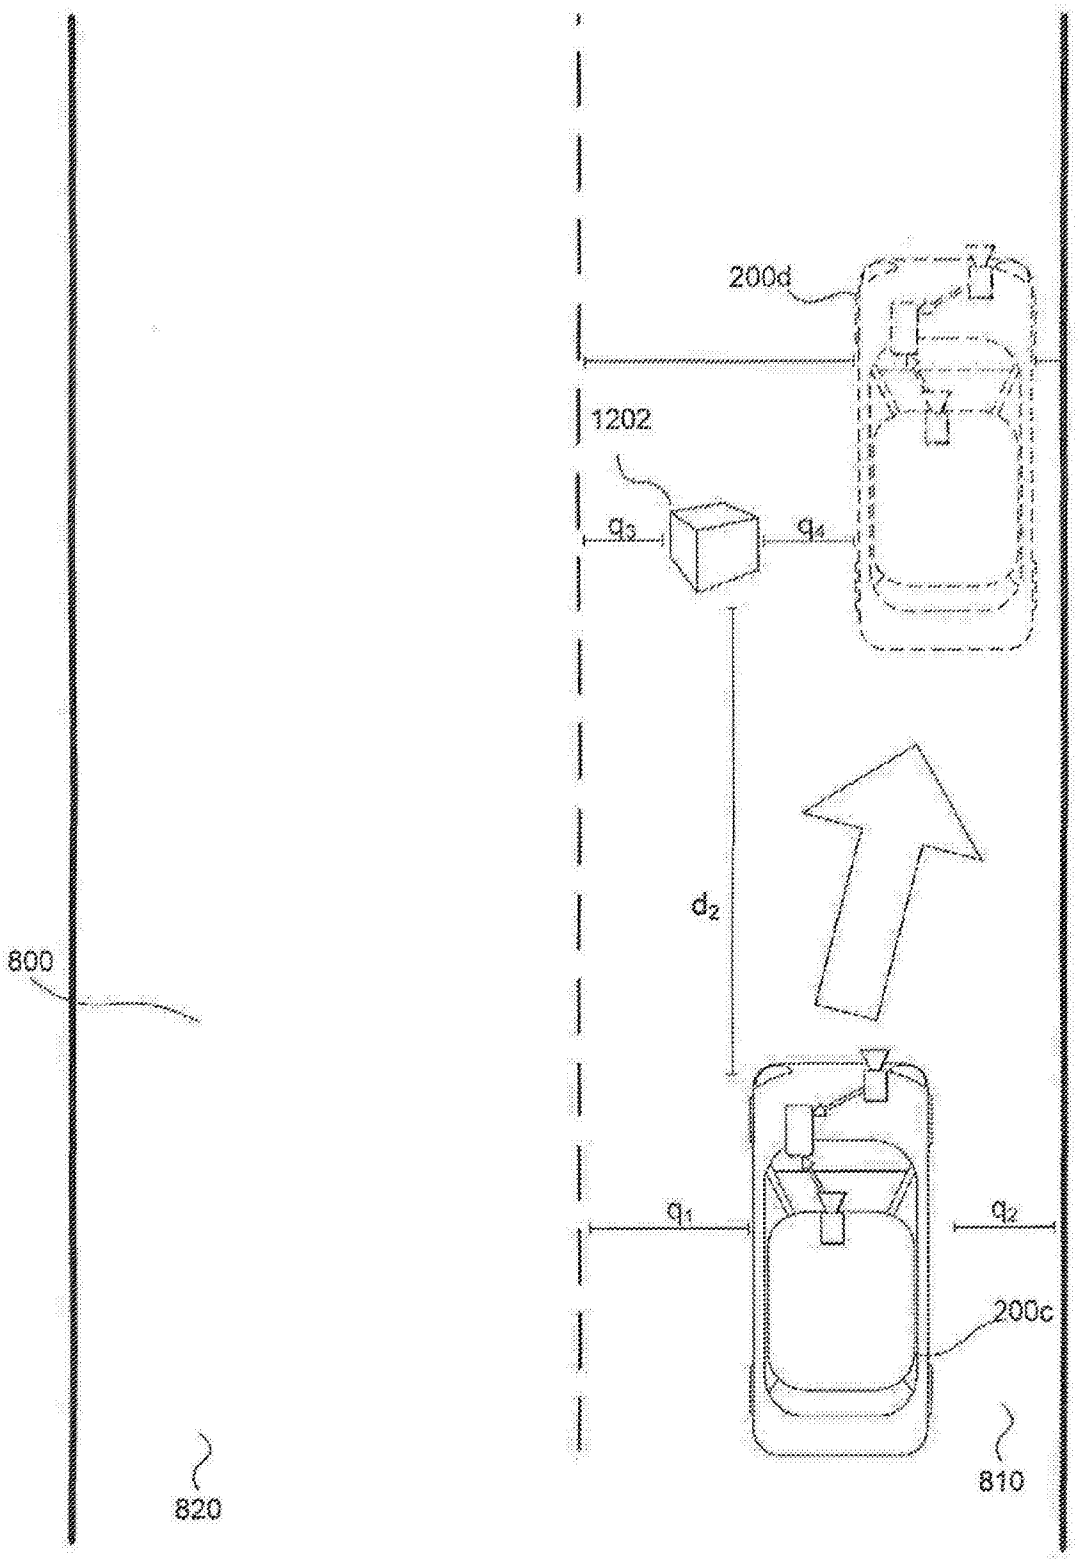


图14B

IPR2025-00943
Tesla EX1002 Page 1016

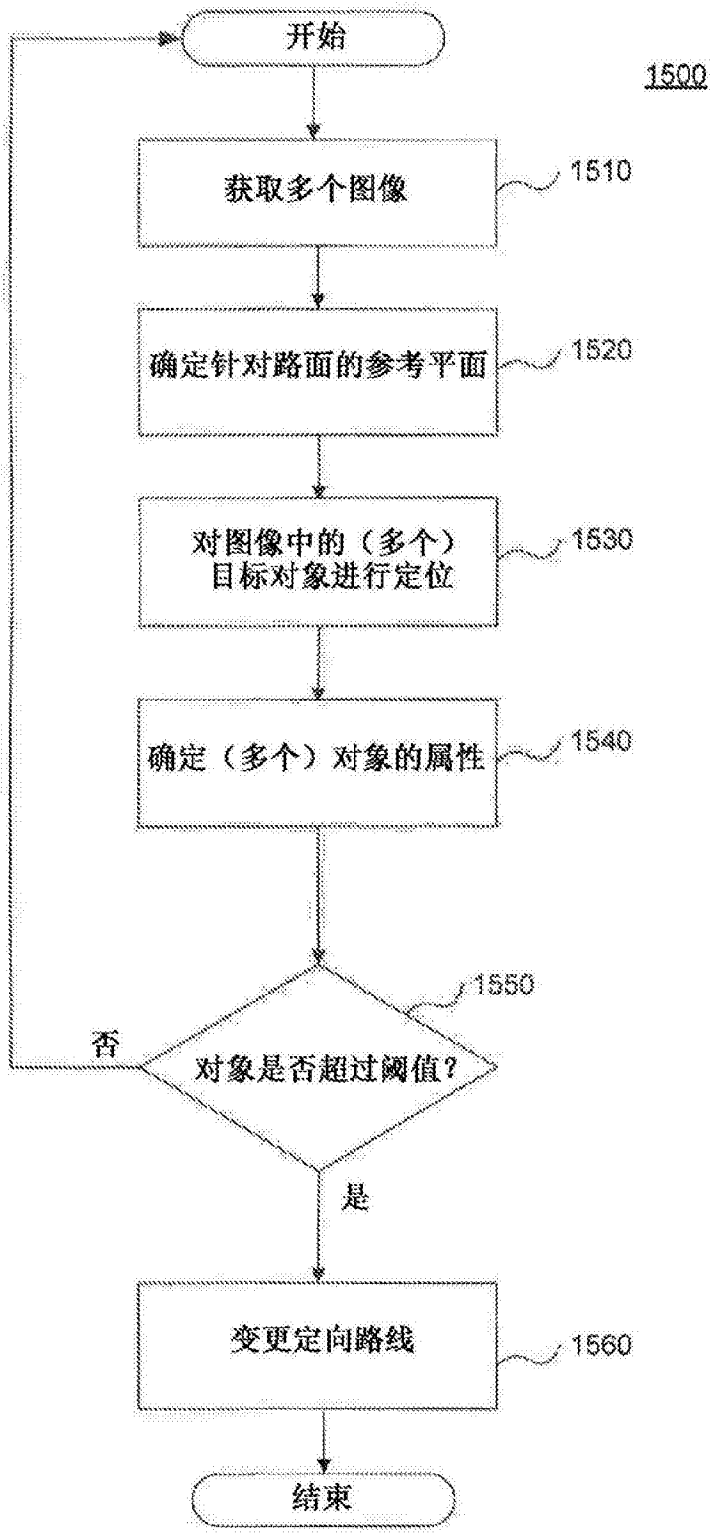


图15

IPR2025-00943
Tesla EX1002 Page 1017

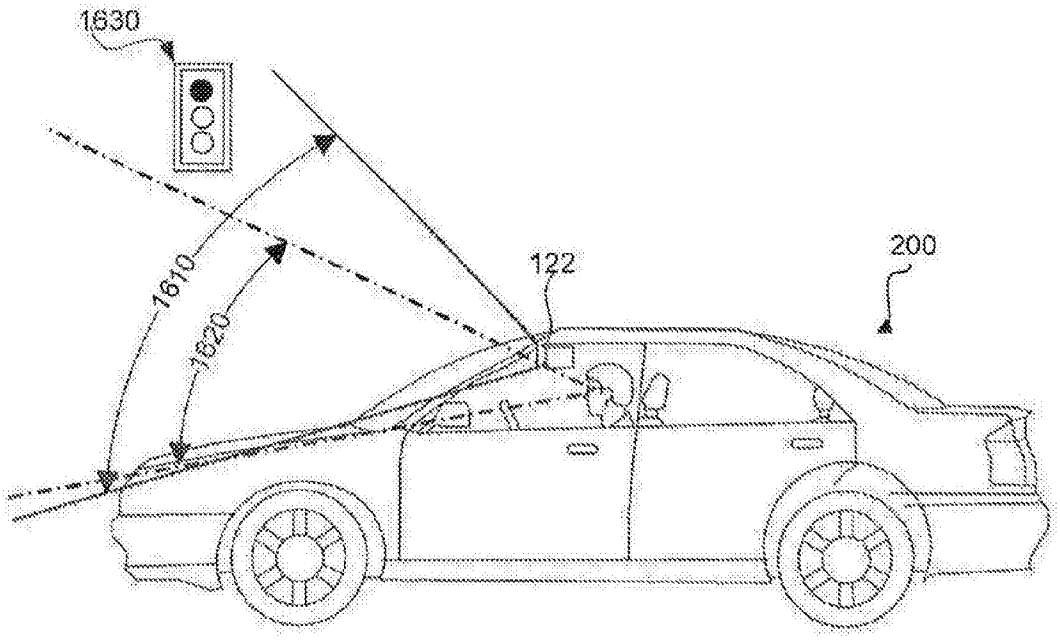


图16

IPR2025-00943
Tesla EX1002 Page 1018

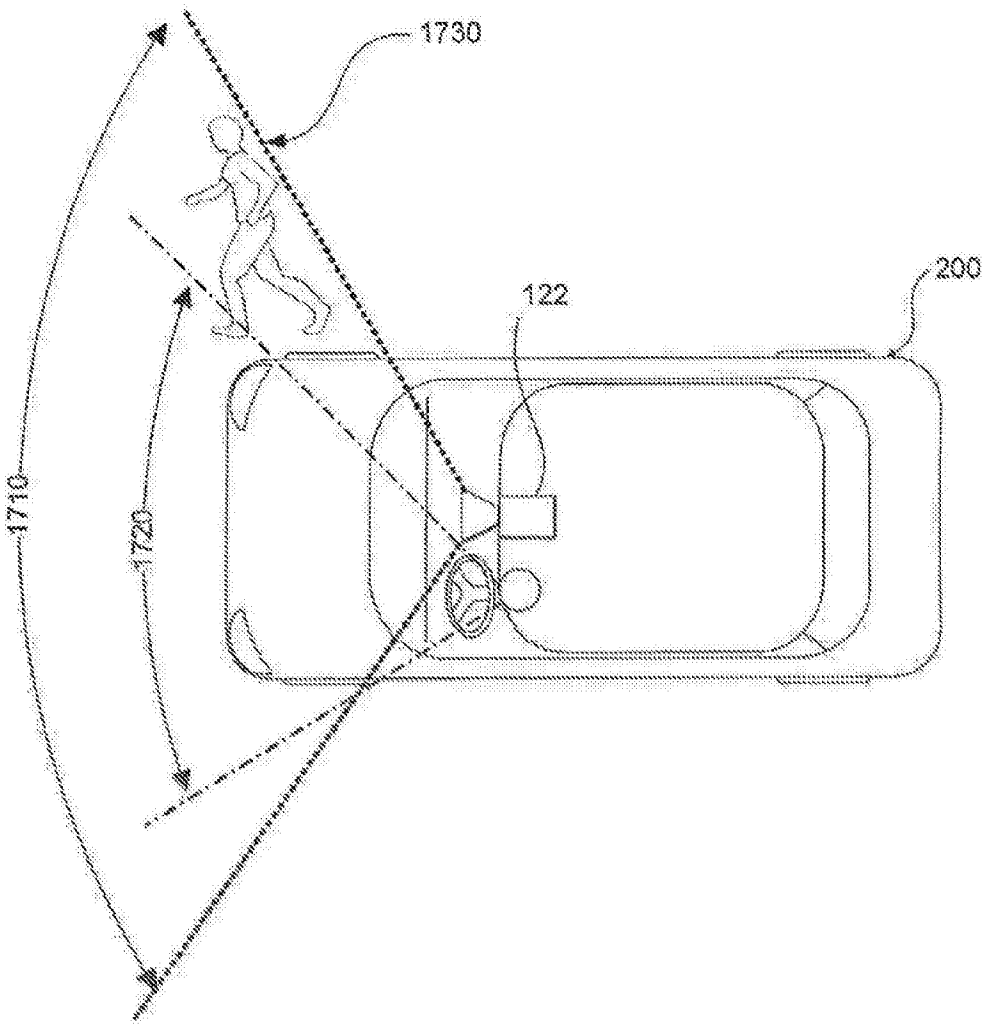


图17

IPR2025-00943
Tesla EX1002 Page 1019

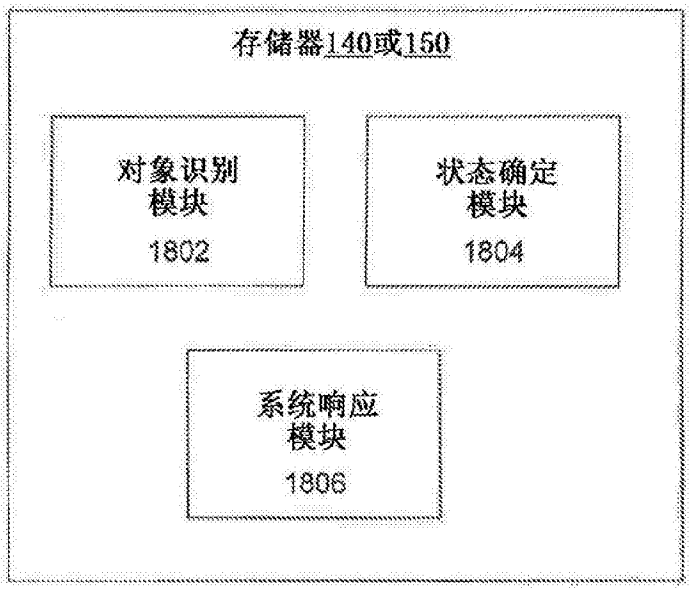


图18

IPR2025-00943
Tesla EX1002 Page 1020

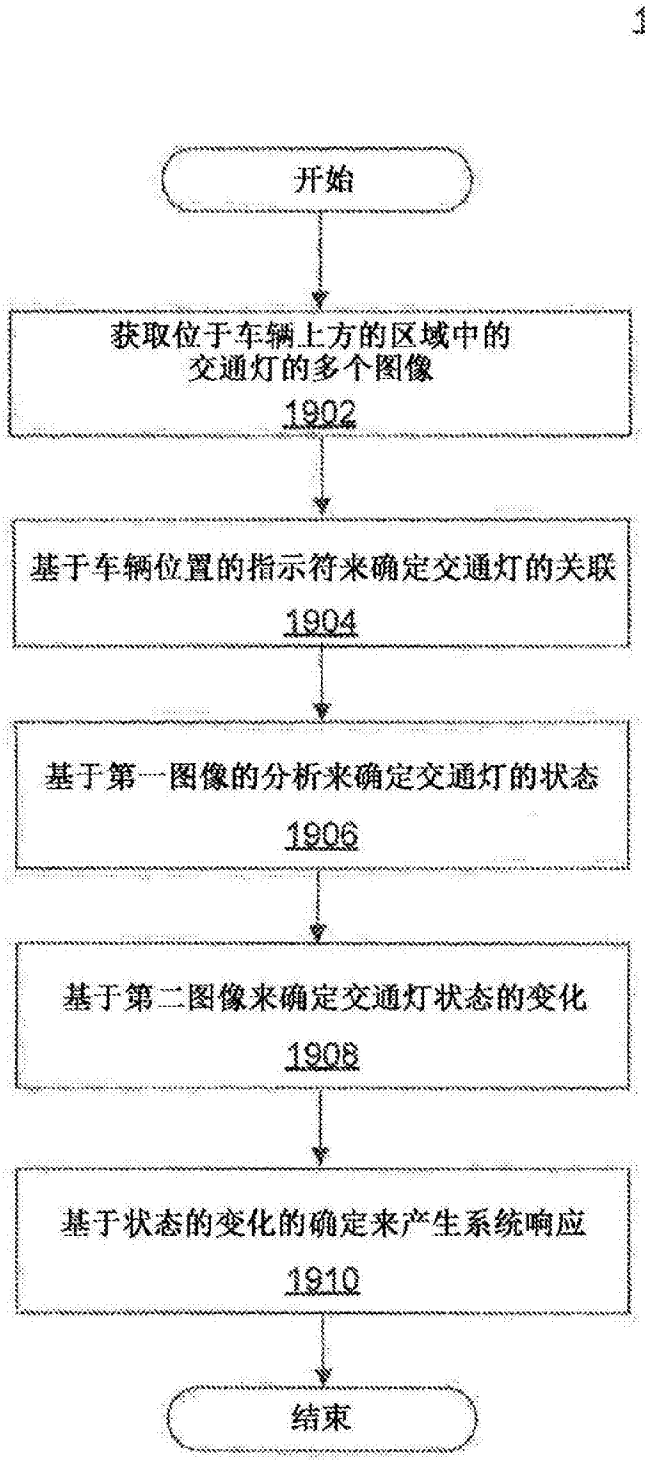


图19

IPR2025-00943
Tesla EX1002 Page 1021

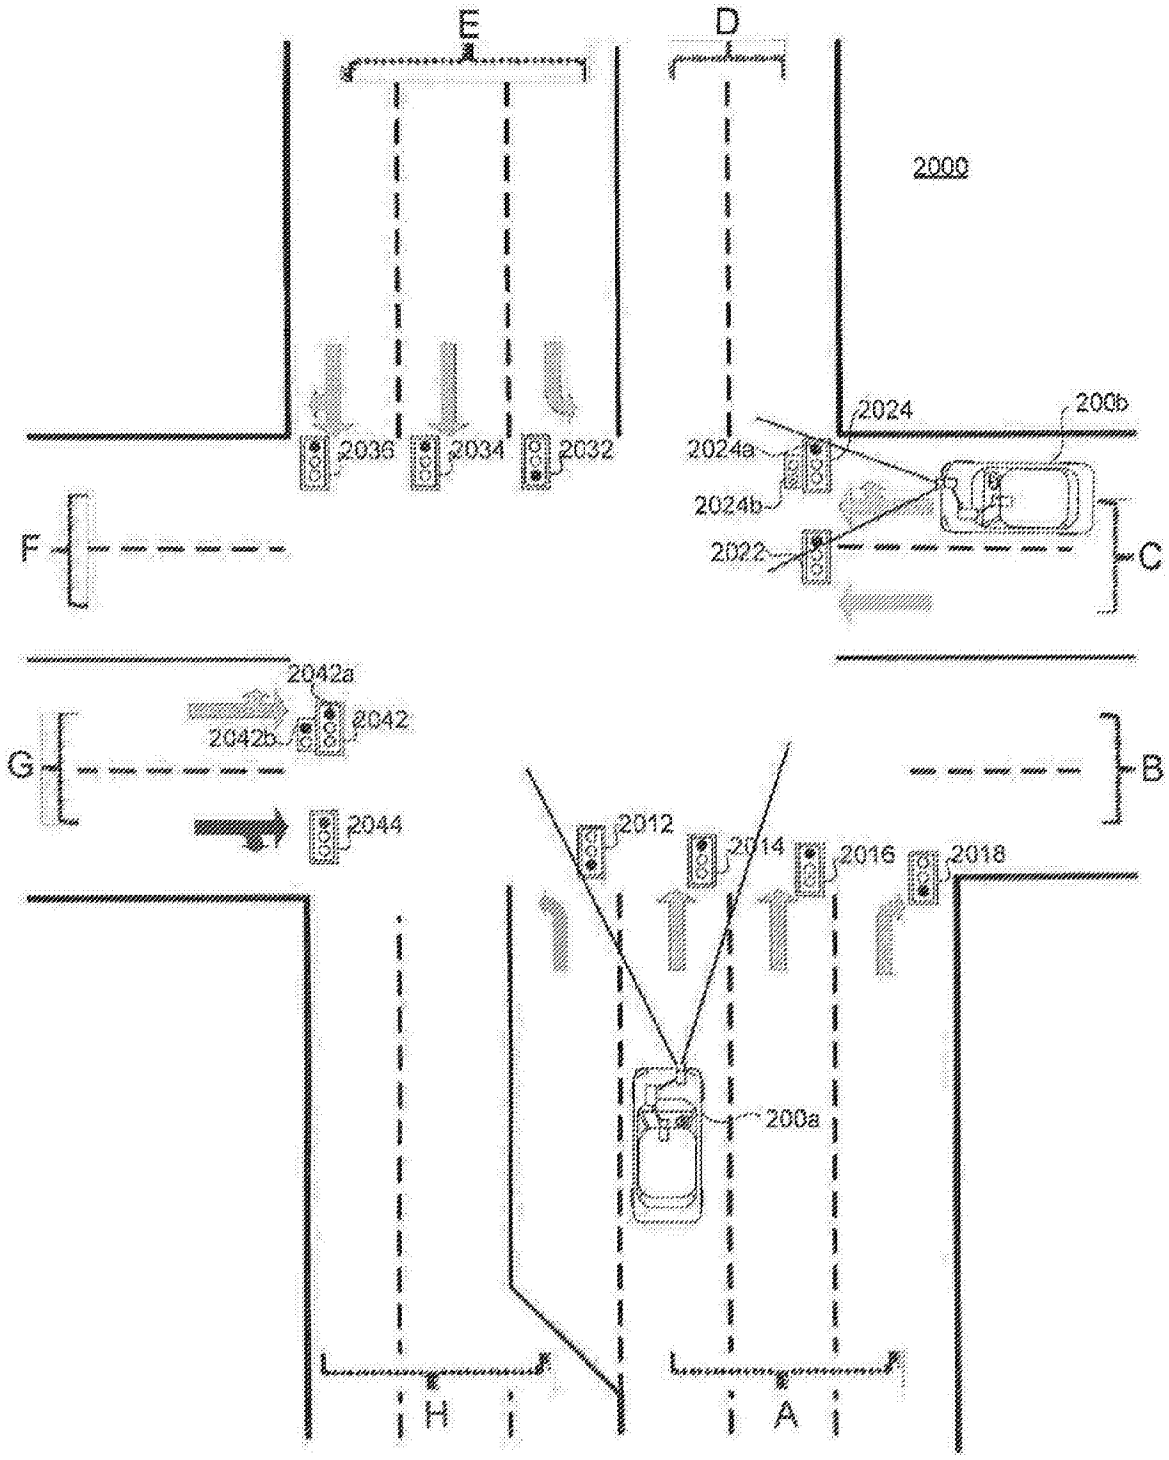


图20

IPR2025-00943
Tesla EX1002 Page 1022

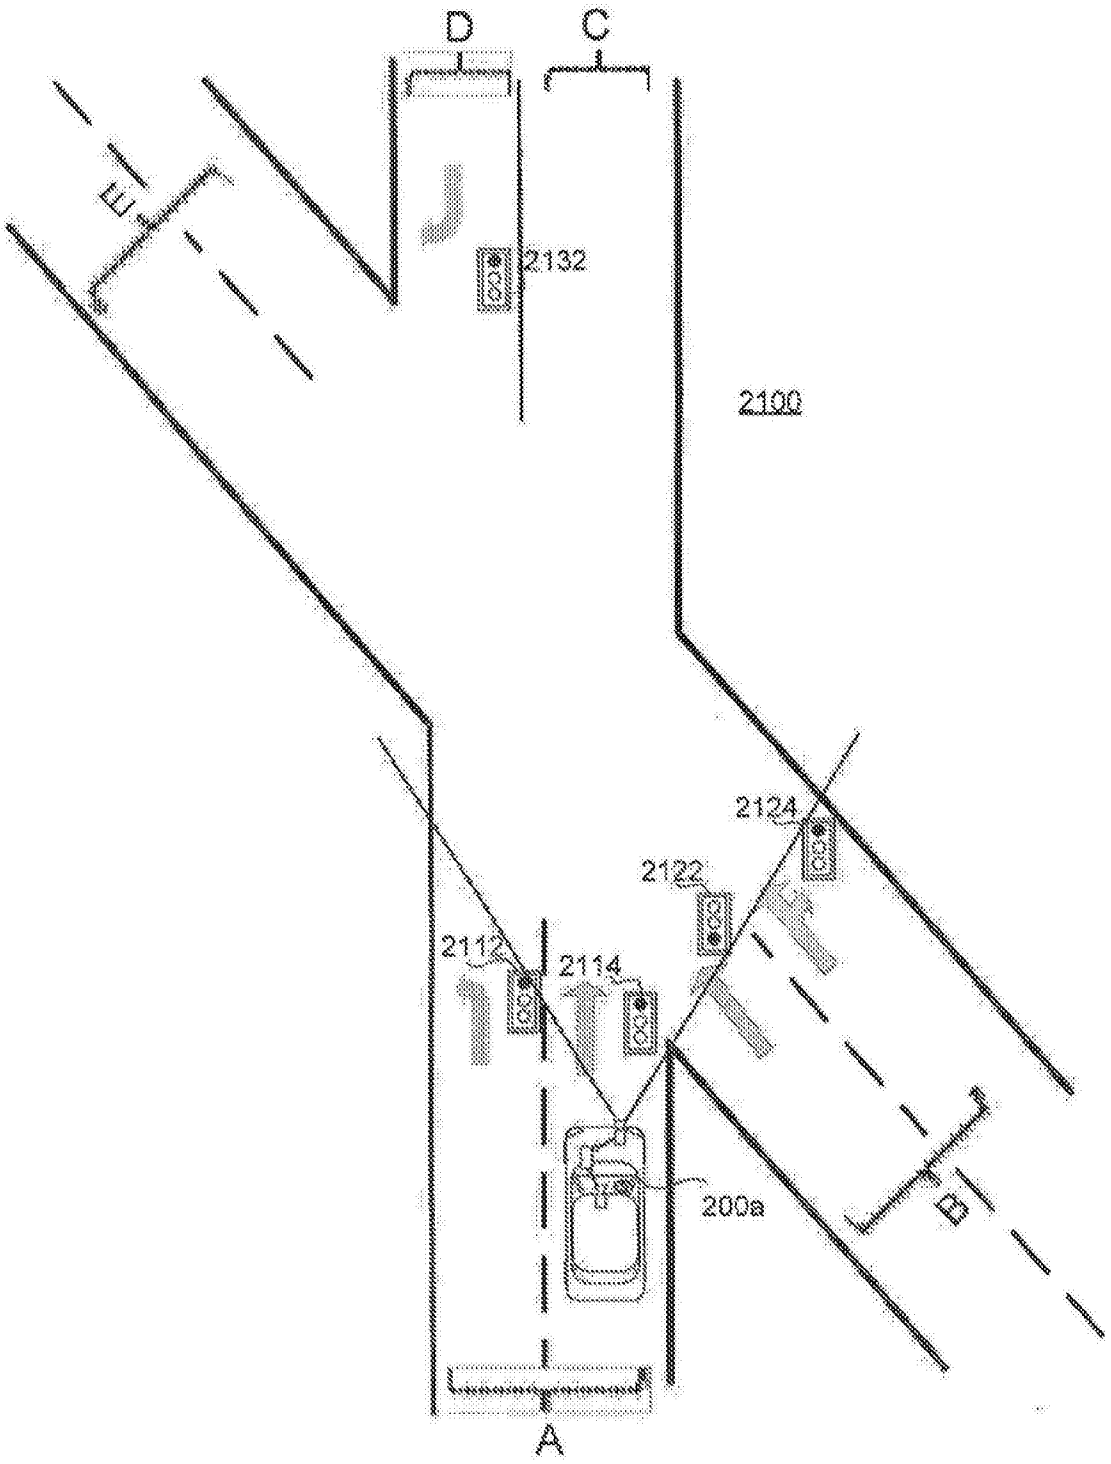


图21

IPR2025-00943
Tesla EX1002 Page 1023

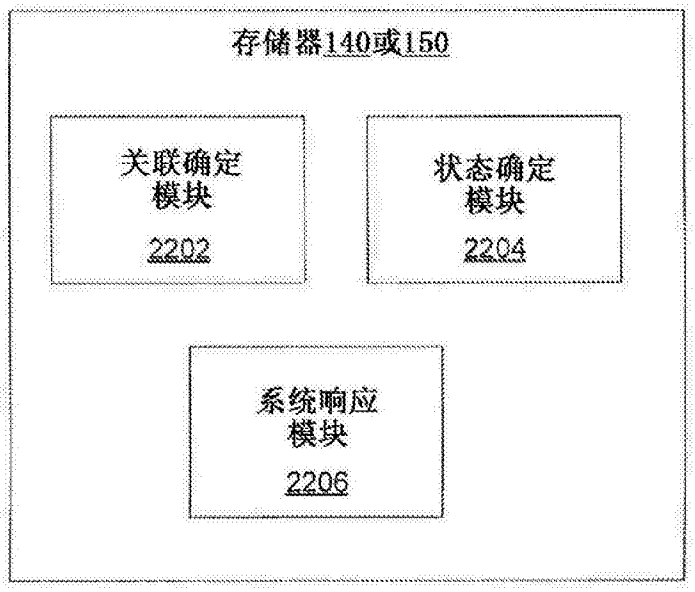


图22

IPR2025-00943
Tesla EX1002 Page 1024

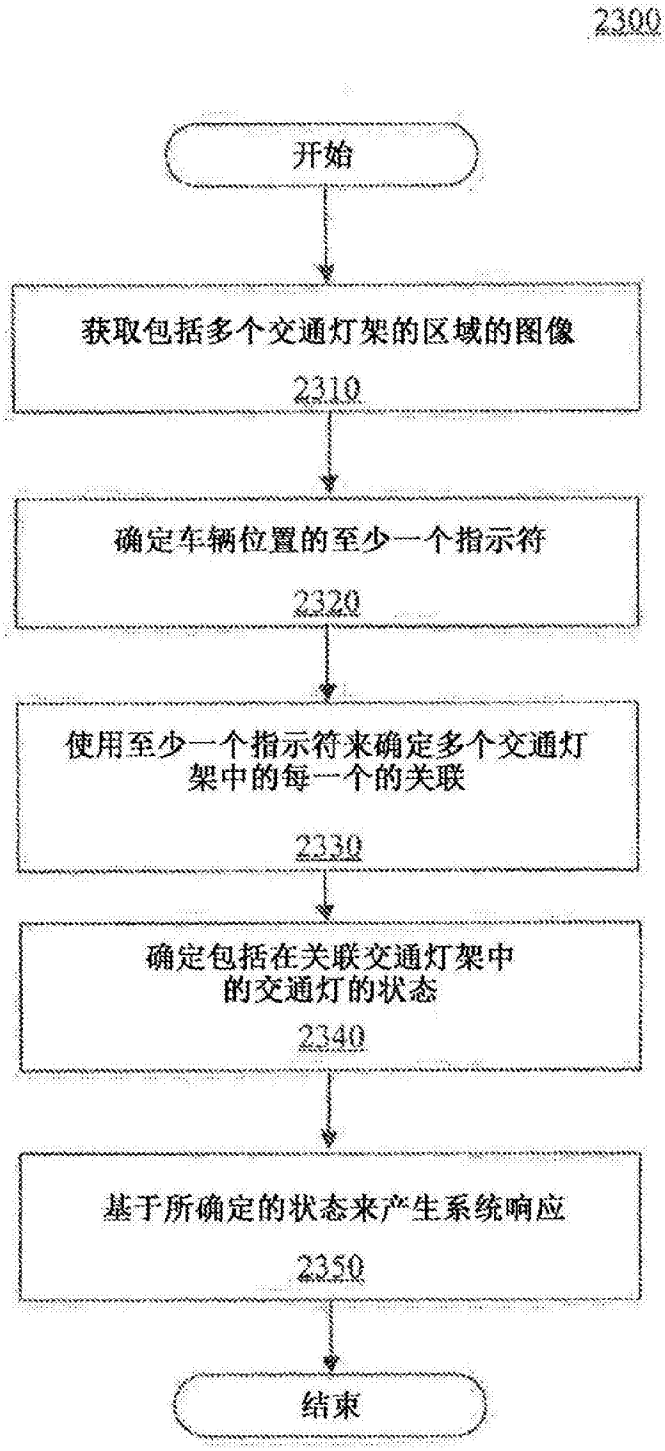


图23

IPR2025-00943
Tesla EX1002 Page 1025

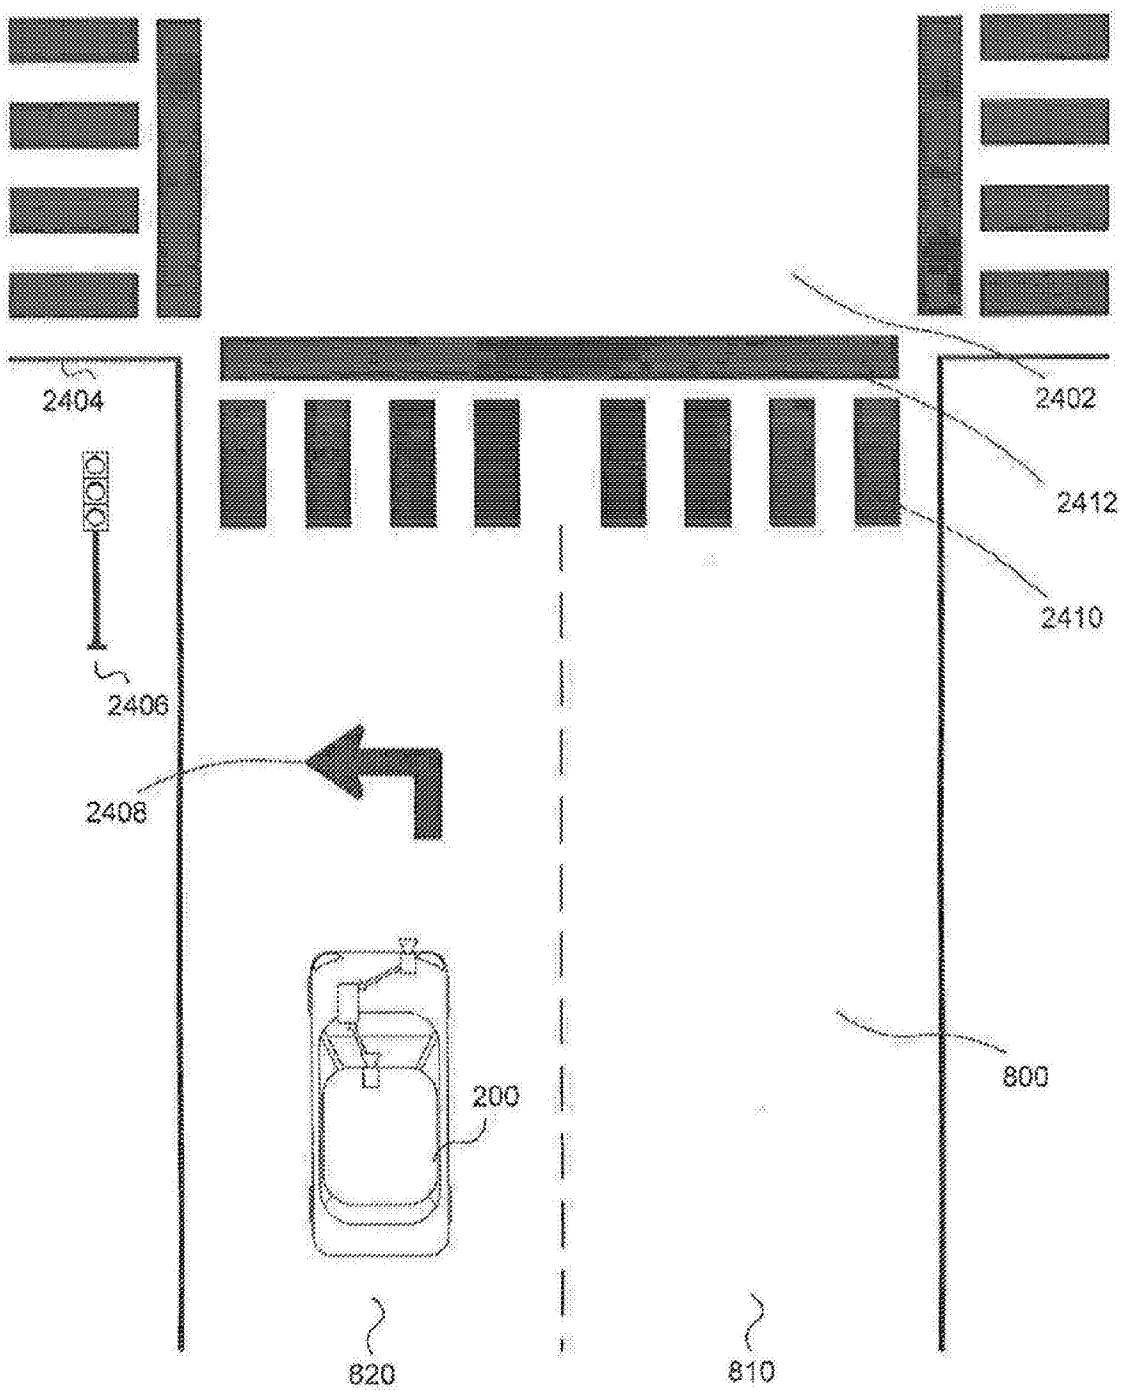


图24

IPR2025-00943
Tesla EX1002 Page 1026

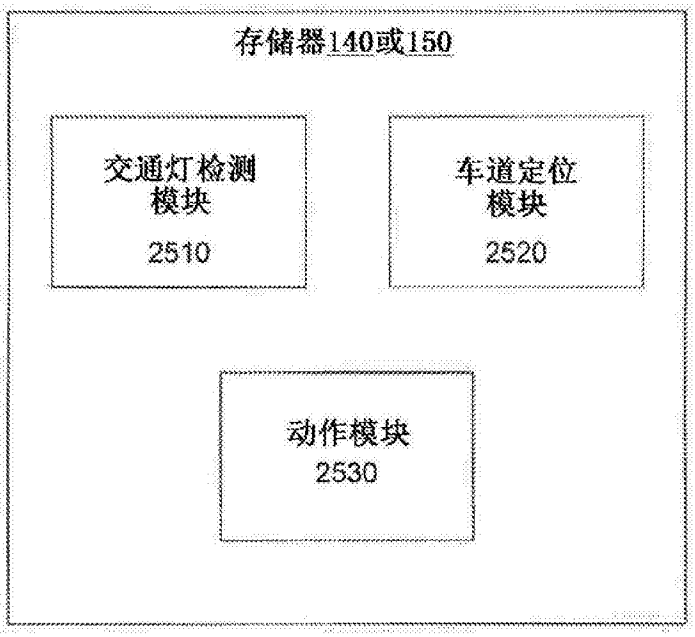


图25

IPR2025-00943
Tesla EX1002 Page 1027

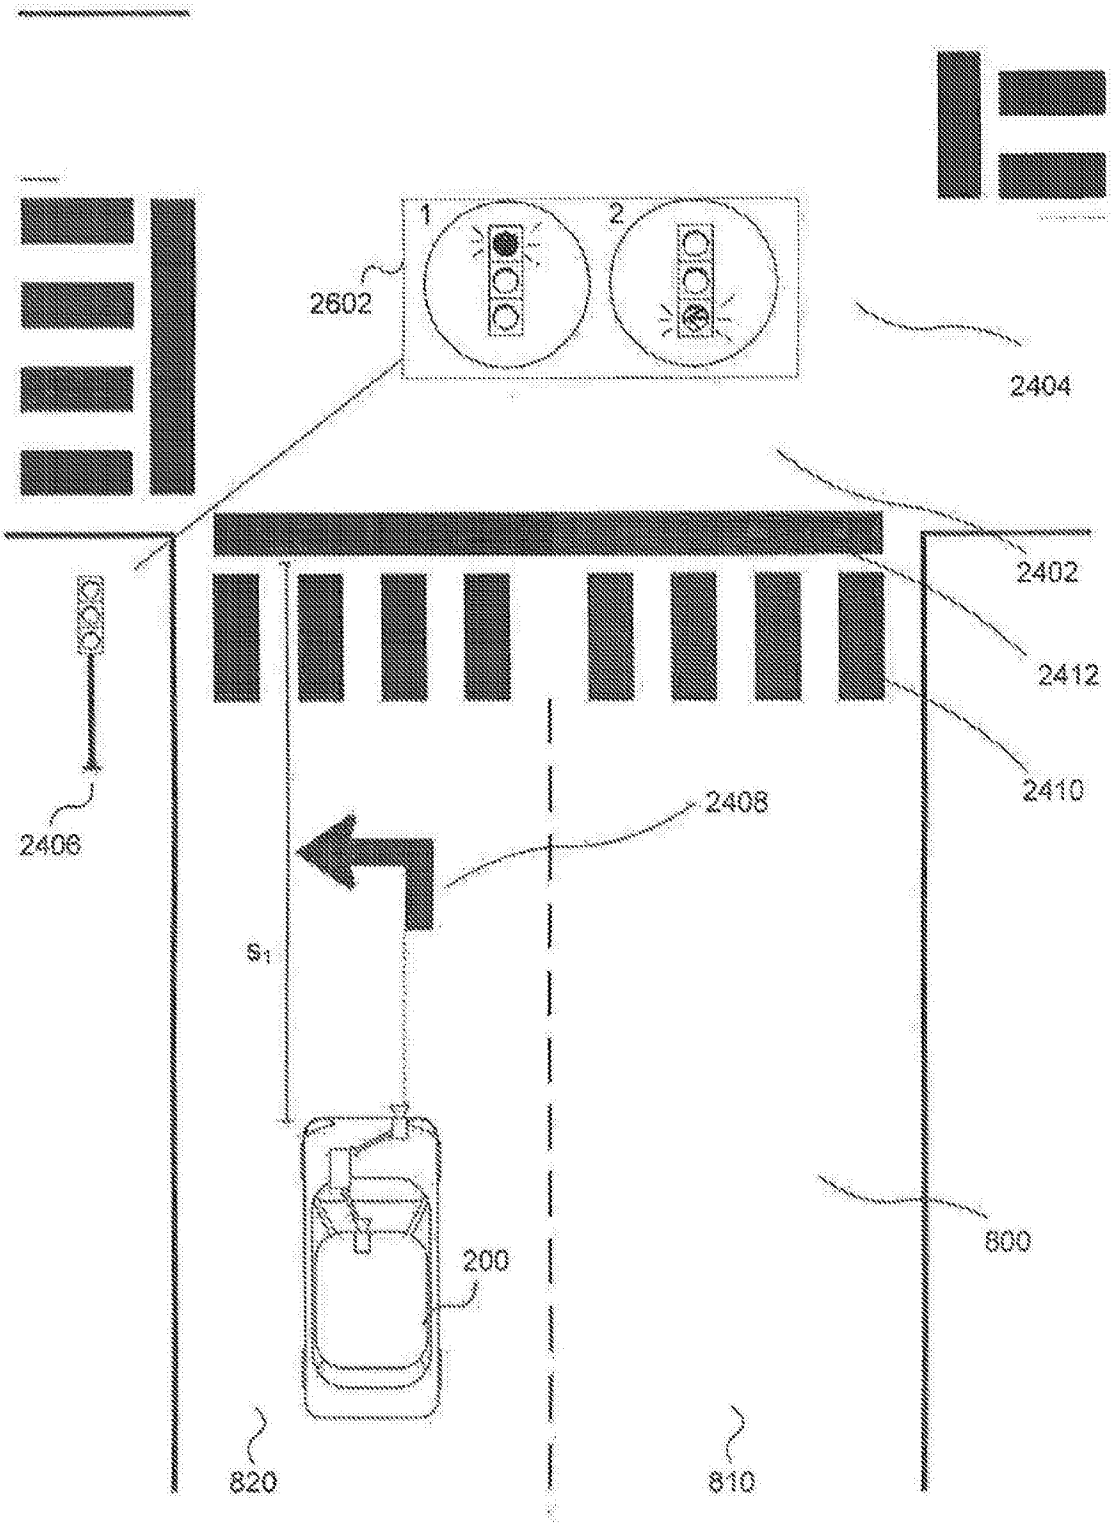


图26A

IPR2025-00943
Tesla EX1002 Page 1028

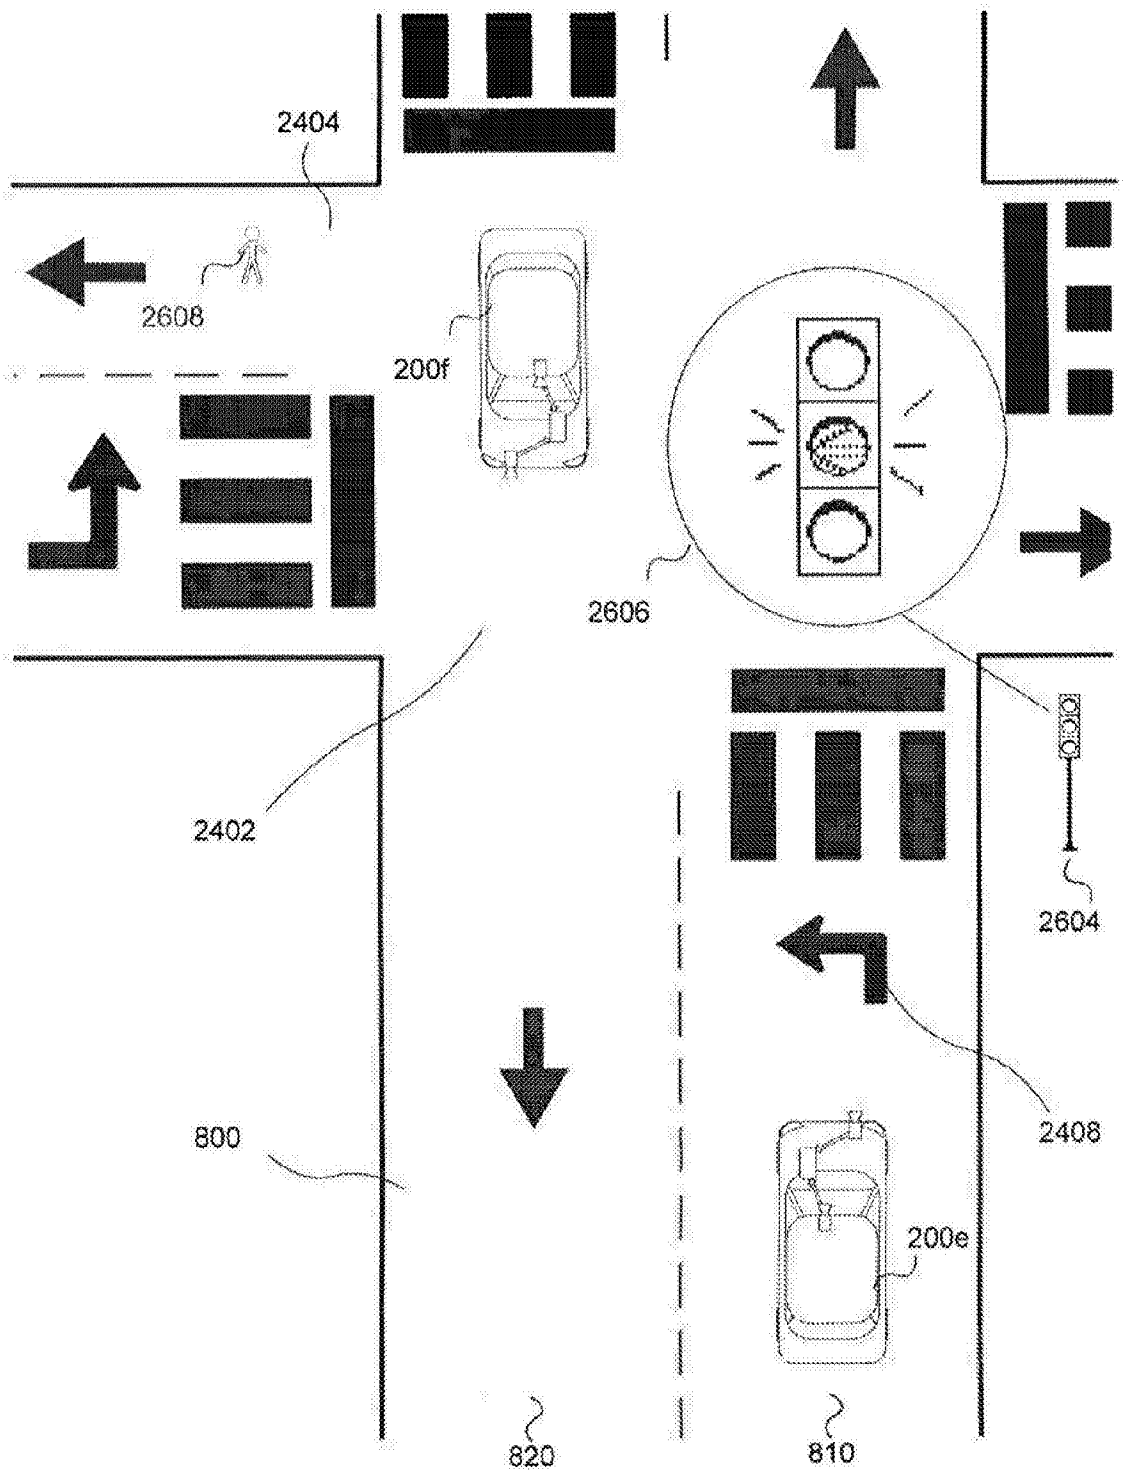


图26B

IPR2025-00943
Tesla EX1002 Page 1029

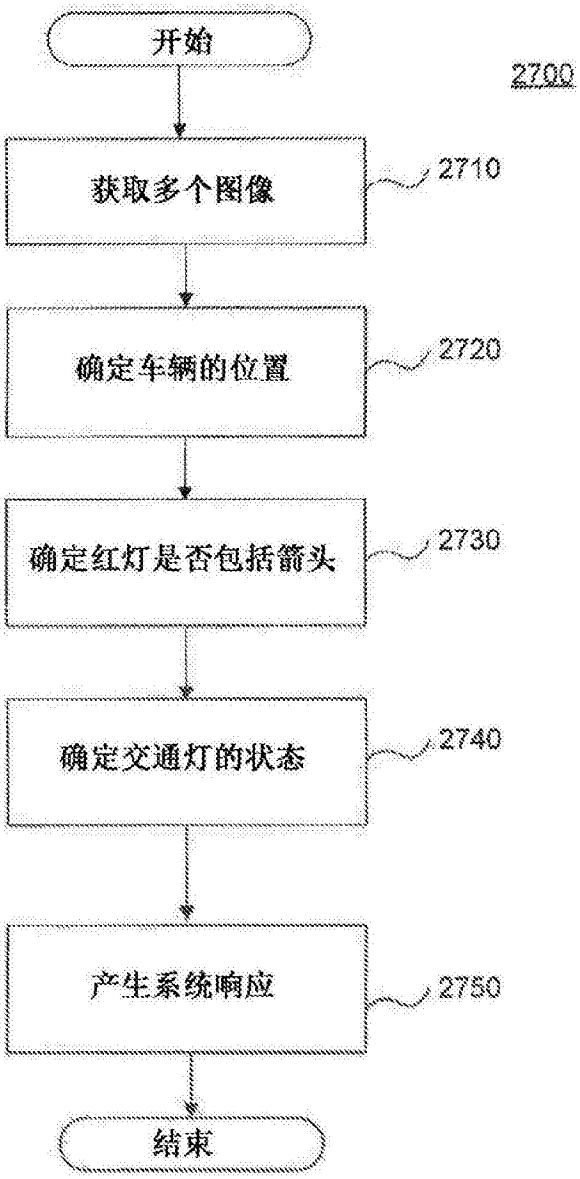


图27

IPR2025-00943
Tesla EX1002 Page 1030

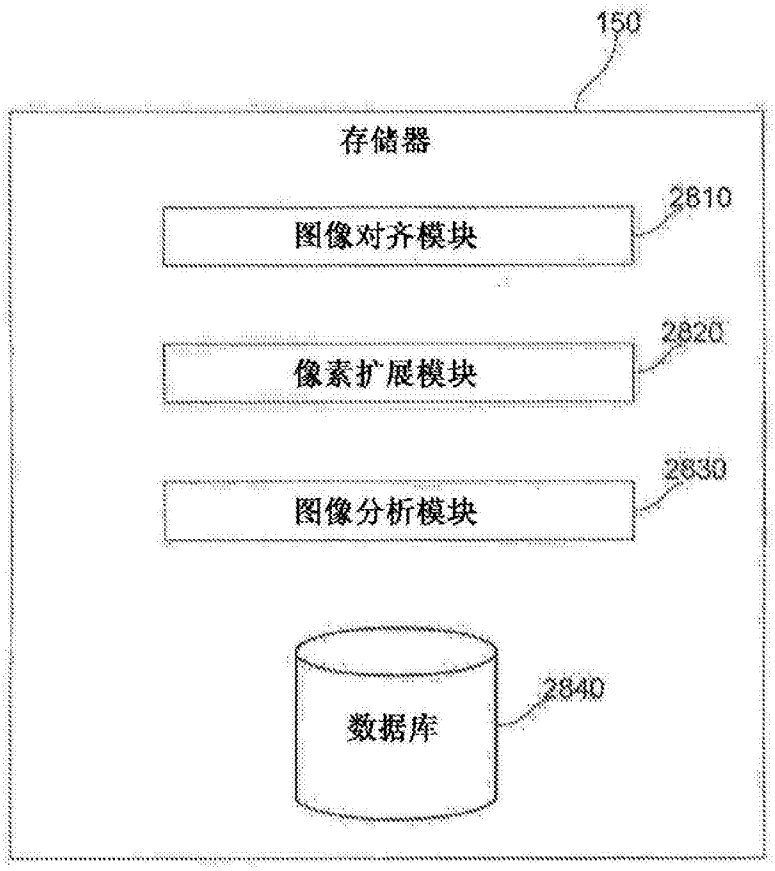


图28

IPR2025-00943
Tesla EX1002 Page 1031

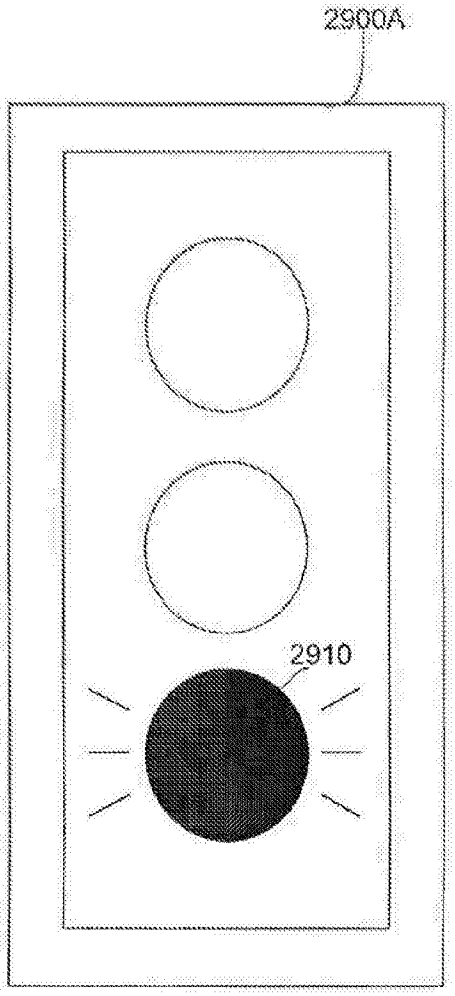


图29A

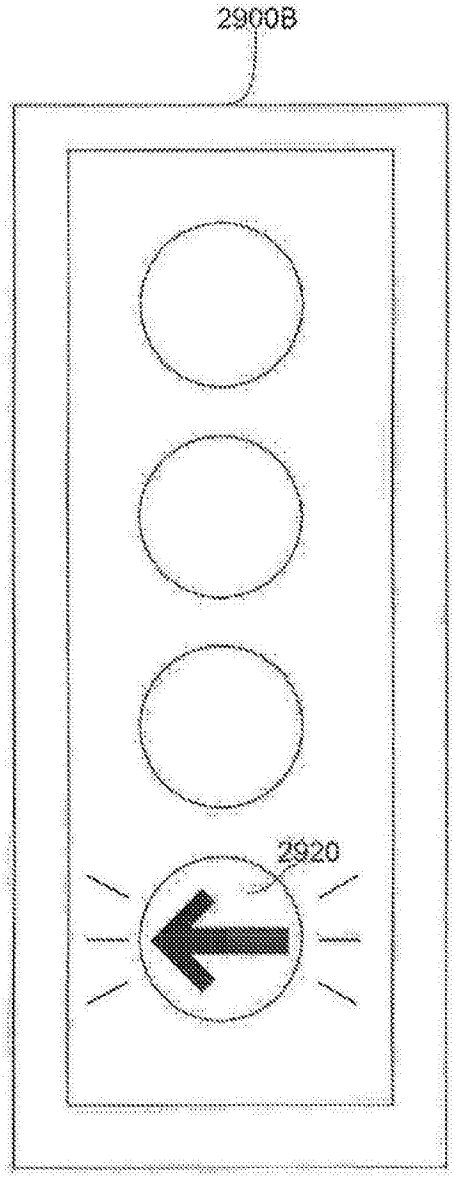


图29B

IPR2025-00943
Tesla EX1002 Page 1032

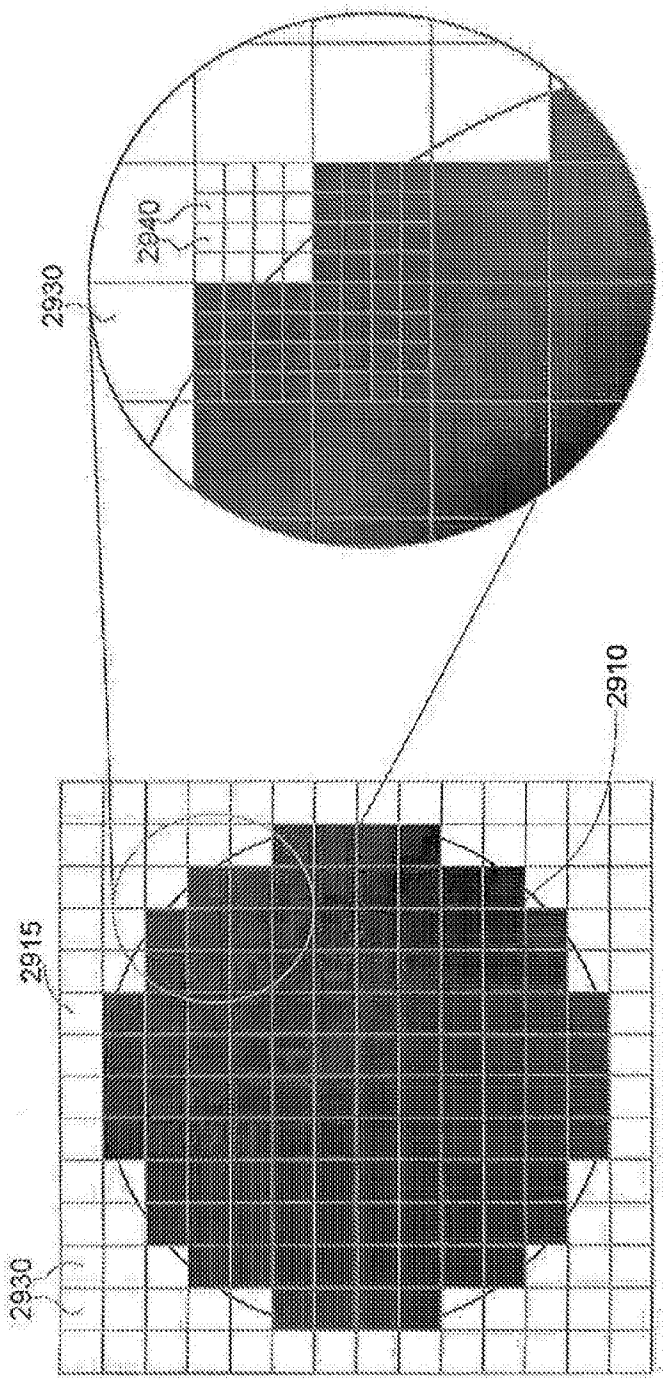


图29C

IPR2025-00943
Tesla EX1002 Page 1033

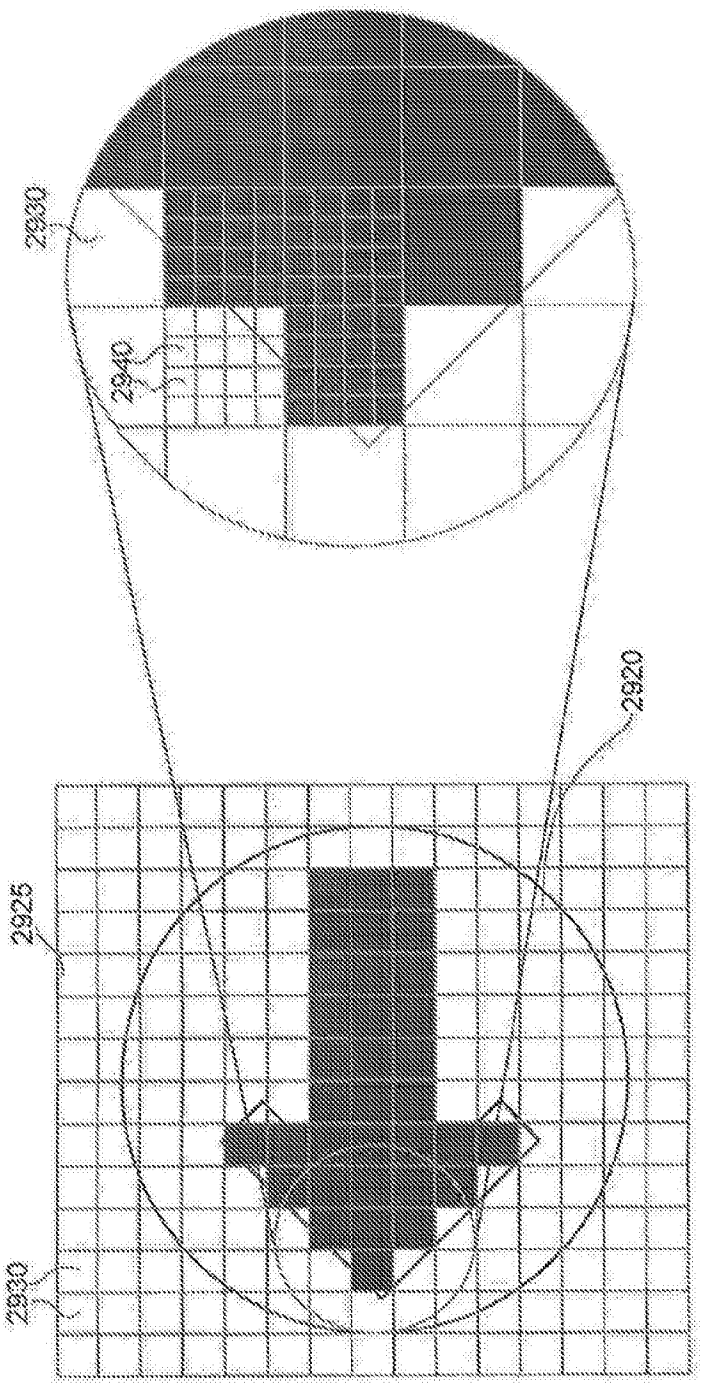


图29D

IPR2025-00943
Tesla EX1002 Page 1034

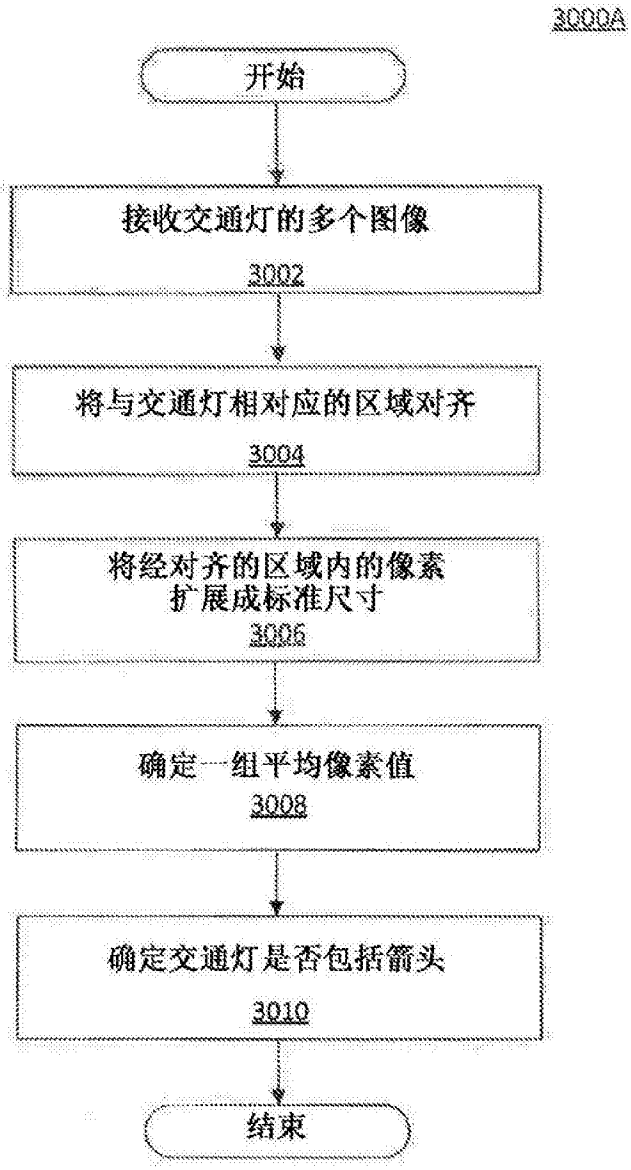


图30A

IPR2025-00943
Tesla EX1002 Page 1035

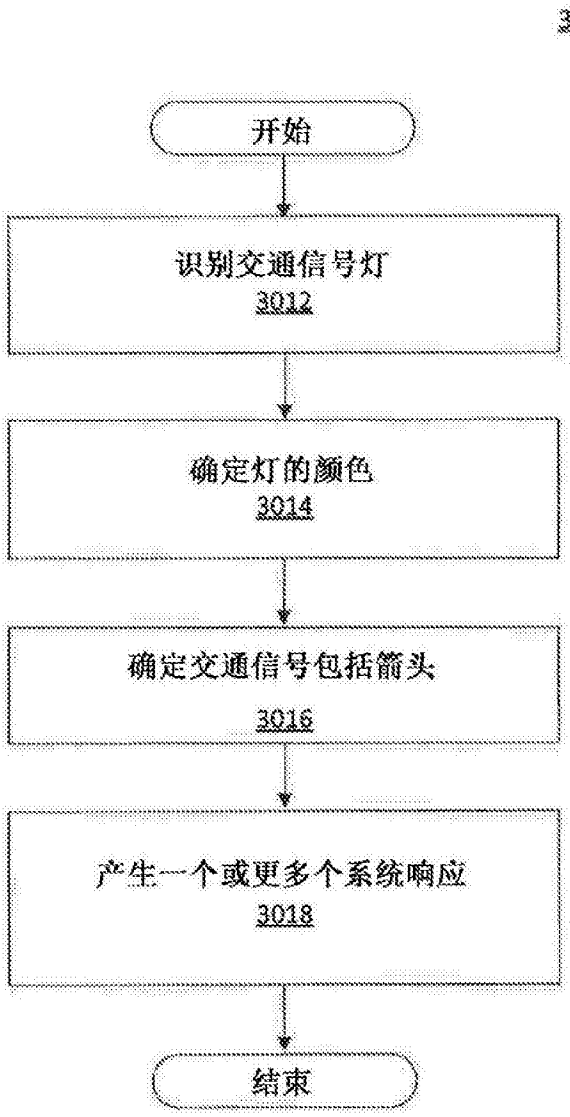


图30B

IPR2025-00943
Tesla EX1002 Page 1036

Espacenet

Bibliographic data: CN107111742 (A) — 2017-08-29

RECOGNITION AND PREDICTION OF LANE CONSTRAINTS AND CONSTRUCTION AREAS IN NAVIGATION

Inventor(s): AHARONY HILA; DYSHEL MARIA; HAREL AMIR; RABANI MEITAL; SHAMAH SHAKED; HORN KOBI ± (希拉·阿哈龙尼, ; 玛丽亚·迪谢尔, ; 阿米尔·哈雷尔, ; 马泽尔·拉巴尼, ; 沙克德·沙玛, ; 科比·霍恩)

Applicant(s): MOBILEYE VISION TECHNOLOGIES LTD ± (无比视视觉技术有限公司)

Classification: **- international:** ***G06K9/00***
- cooperative: **B60W10/04 (US)**; **B60W10/18 (US)**; **B60W10/20 (US)**; **B60W30/0953 (US)**; **B60W30/0956 (US)**; **G06K9/00798 (EP, CN, US)**; **G08G1/09623 (EP, US)**; B60W2420/403 (US); B60W2710/18 (US); B60W2710/20 (US); B60W2720/10 (US)

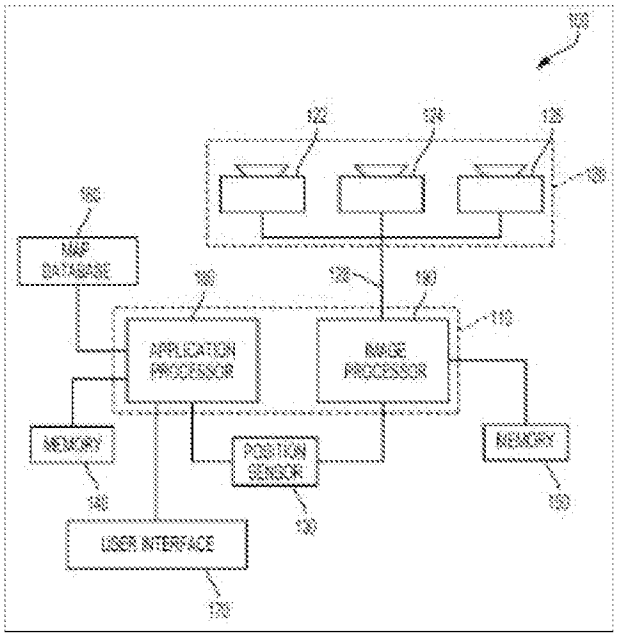
Application number: CN20158055136 20150817 Global Dossier

Priority number(s): US201462038426P 20140818 ; WO2015IL50833 20150817

Also published as: CN107111742 (B) CN113158820 (A) EP3183688 (A1) US10150473 (B2) US10926763 (B2) more

Abstract of CN107111742 (A)

Systems and methods use cameras to provide autonomous and/or driver-assist navigation features. In some implementations, techniques for predicting the location of first roadway lane constraints are provided. The system may receive multiple images of a roadway in a vicinity of a vehicle, recognize a first roadway lane constraint, and, when lane prediction conditions are determined to be satisfied, predict a location of a second roadway lane constraint. In some implementations, techniques for detecting and responding to construction zones are provided. The system may receive multiple images of a roadway in a vicinity of a vehicle, recognize indicators of a construction zone in the images, determine that the vehicle is proximate to a construction zone, and output a signal indicating that the vehicle is proximate to a construction zone.



IPR2025-00943
Tesla EX1002 Page 1038

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

DESCRIPTION CN107111742A

Recognition and prediction of lane restrictions and construction areas in navigation

[0001]

Cross-references to related applications

[0002]

This application claims the benefit of U.S. Provisional Patent Application No. 62/038,426, the disclosure of which is incorporated herein by reference.

[0003]

Public domain

[0004]

The present disclosure generally relates to autonomous driving and/or driver assistance technology, and more particularly to systems and methods that use cameras to provide features of autonomous driving and/or driver assistance technology.

[0005]

background

[0006]

With the continuous advancement of technology, the goal of a fully autonomous vehicle capable of navigating on the road will soon be achieved.

First, an autonomous vehicle can recognize its environment and navigate without input from a human operator.

Autonomous vehicles can also consider various factors and make appropriate decisions based on these factors to safely and accurately reach the desired destination.

For example, when a vehicle usually travels on a road, various objects such as other vehicles and pedestrians are encountered.

The autonomous driving system can recognize these objects in the vehicle's environment and take appropriate and timely actions to avoid collisions.

In addition, the autonomous driving system can recognize other indicators that regulate the movement of the vehicle (for example, when the vehicle must stop and can go, the speed that the vehicle must not exceed, where the vehicle must be located on the road, etc.), such as traffic signals, traffic signs and signs And lane markings. Autonomous driving systems may need to determine when the vehicle should change lanes, turn at intersections, change roads, etc. It is obvious from these examples that in order to provide an autonomous vehicle capable of safe and accurate navigation, many factors may need to be addressed.

[0007]

Overview

[0008]

Current autonomous or driver-assisted (e.g., partially autonomous) vehicle systems may not be able to fully recognize lanes when lane lines (e.g., lane markings) or other markings do not clearly indicate the location of lane restrictions on the road.

For example, the current system may not work properly when the lane restriction cannot be recognized from the line or mark detected in the received road image; or the current system may not work when the lane restriction does not exist in reality. Accurately identify lane restrictions.

Current autonomous or driver-assisted vehicle systems may also not be able to adequately recognize or respond to construction areas on the road, and may require modified vehicle control in construction areas, and may respond to certain road elements in different ways in construction areas. Is advantageous. The inventor has recognized these and other problems, and these and other problems will become apparent through the present disclosure and solved by the technology disclosed herein.

[0009]

Some embodiments consistent with the present disclosure provide systems and methods for autonomous or driver-assisted vehicle navigation.

Some disclosed embodiments may use cameras to provide autonomous vehicle navigation features. For example, a disclosed system consistent with some disclosed embodiments may include one, two, or more cameras that monitor the environment of the vehicle and cause a navigation response based on the analysis of the images captured by the one or more cameras.

[0010]

According to some embodiments, the computer system includes a memory storing instructions and a processor that executes the instructions to enable the system to perform the following operations: receive multiple images of the road near the moving vehicle from the camera; identify the first road lane restriction based on the multiple images ; Based on multiple images, determine that one or more defined lane prediction conditions are met; based on the determination that one or more defined lane prediction conditions are satisfied, predict the position of the lane restriction on the second road; and enable moving vehicles to avoid the first Lane restriction on the first road and lane restriction on the second road.

[0011]

According to some embodiments, the computer system includes a memory that stores instructions and a processor that executes the instructions to enable the system to perform the following operations: receive multiple images of roads near the moving vehicle from the camera; identify one of the construction areas from the multiple images Or multiple indicators; determining that the vehicle is approaching the construction area based on one or more indicators; and, according to the determination that the vehicle is approaching the construction area, output a signal indicating that the vehicle is approaching the construction area.

[0012]

Brief description of the drawings

[0013]

Figure 1 is a graphical representation of an exemplary system consistent with the disclosed embodiments.

[0014]

Figure 2A is a graphical side view representation of an exemplary vehicle including a system consistent with the disclosed embodiments.

[0015]

Figure 2B is a graphical top view representation of the vehicle and system shown in Figure 2A consistent with the disclosed embodiments.

[0016]

Figure 2C is a graphical top view representation of another vehicle embodiment including a system consistent with the disclosed embodiment.

[0017]

Figure 2D is a graphical top view representation of yet another vehicle embodiment including a system consistent with the disclosed embodiment.

[0018]

Figure 2E is a graphical top view representation of yet another vehicle embodiment including a system consistent with the disclosed embodiment.

[0019]

Figure 2F is a graphical representation of an exemplary vehicle control system consistent with the disclosed embodiments.

[0020]

Figure 3A is a graphical representation of the interior of a vehicle including a rear view mirror and a user interface for the vehicle imaging system consistent with the disclosed embodiment.

[0021]

FIG. 3B is an illustration of an example of a camera base configured to be positioned behind a rearview mirror and facing the windshield of a vehicle consistent with the disclosed embodiment.

[0022]

Fig. 3C is an illustration of the camera base shown in Fig. 3B from a different angle consistent with the disclosed embodiment.

[0023]

3D is an illustration of an example of a camera base configured to be positioned behind the rearview mirror and facing the windshield of the vehicle consistent with the disclosed embodiment.

[0024]

Figure 4 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

[0025]

FIG. 5A is a flowchart showing an exemplary process for eliciting one or more navigation responses based on monocular image analysis consistent with the disclosed embodiments.

[0026]

FIG. 5B is a flowchart showing an exemplary process for detecting one or more vehicles and/or pedestrians in a set of images consistent with the disclosed embodiments.

[0027]

FIG. 5C is a flowchart showing an exemplary process for detecting road markings and/or lane geometry information in a set of images consistent with the disclosed embodiments.

[0028]

FIG. 5D is a flowchart showing an exemplary process for detecting traffic lights in a set of images consistent with the disclosed embodiments.

[0029]

FIG. 5E is a flowchart showing an exemplary process for eliciting one or more navigation responses based on a vehicle path consistent with the disclosed embodiment.

[0030]

FIG. 5F is a flowchart showing an exemplary process for determining whether a leading vehicle is changing lanes, consistent with the disclosed embodiment.

[0031]

FIG. 6 is a flowchart showing an exemplary process for eliciting one or more navigation responses based on stereoscopic image analysis consistent with the disclosed embodiments.

[0032]

FIG. 7 is a flowchart showing an exemplary process for causing one or more navigation responses based on analysis of three sets of images consistent with the disclosed embodiment.

[0033]

Figures 8A-8C are depictions of vehicles traveling along a road according to some embodiments.

[0034]

Figures 9A-9C are depictions of vehicles traveling along a road according to some embodiments.

[0035]

Figure 10A is a depiction from a perspective view of a camera associated with a system according to some embodiments.

[0036]

Figure 10B is a depiction of a vehicle traveling along a road according to some embodiments.

[0037]

Figure 11 is a graphical representation of the memory of an exemplary system consistent with some disclosed embodiments.

[0038]

Figure 12 is a flowchart of an exemplary process according to some embodiments.

[0039]

Figure 13 is a depiction of a vehicle traveling along a road according to some embodiments.

[0040]

Figure 14 is a graphical representation of the memory of an exemplary system consistent with some disclosed embodiments.

[0041]

Figure 15 is a flowchart of an exemplary process according to some embodiments.

[0042]

A detailed description

[0043]

The following detailed description refers to the accompanying drawings.

Whenever possible, the same reference numbers are used in the drawings and the following description to refer to the same or similar parts.

Although several illustrative embodiments have been described herein, modifications, adaptations, and other implementations are also possible.

For example, the components illustrated in the drawings may be replaced, added, or modified, and the illustrative methods described herein may be modified by replacing, rearranging, removing, or adding steps to the disclosed method.

Therefore, the following detailed description is not limited to the disclosed embodiments and examples.

[0044]

The disclosed embodiments provide systems and methods that use cameras to provide features of autonomous navigation, vehicle control, and/or driver assistance technology.

Driver assistance technology, rather than fully autonomous driving, refers to any suitable technology used to assist the driver in the navigation and control of their vehicle, such as LKA (Lane Keep Assist), LDW (Lane Departure Warning), vehicle acceleration, The deceleration of the vehicle, the steering of the vehicle, controlling the braking of the vehicle, preparing the vehicle for an inevitable collision, and/or controlling any aspect of vehicle operation or any other suitable way to assist the driver in controlling any such aspect.

In various embodiments, the system may include one, two, or more cameras that monitor the environment of the vehicle.

In some embodiments, the system may provide techniques for predicting the location of road lane restrictions and/or techniques for detecting and responding to construction areas.

[0045]

Figure 1 is a block diagram representation of a system 100 consistent with an exemplary disclosed embodiment.

The system 100 may include various components depending on the requirements of a particular implementation.

In some embodiments, the system 100 may include a processing unit 110, an image acquisition unit 120, a positioning sensor 130, one or more memory units 140, 150, a map database 160, and a user interface 170.

IPR2025-00943
Tesla EX1002 Page 1043

The processing unit 110 may include one or more processing devices.

In some embodiments, the processing unit 110 may include an application processor 180, an image processor 190, or any other suitable processing device.

Similarly, the image capture unit 120 may include any number of image capture devices and components depending on the requirements of a particular application.

In some embodiments, the image capture unit 120 may include one or more image capture devices (for example, cameras), such as the image capture device 122, the image capture device 124, and the image capture device 126.

The system 100 may also include a data interface 128 that communicatively connects the processing unit 110 to the image capture device 120.

For example, the data interface 128 may include any wired and/or wireless link for transmitting the image data collected by the image acquisition device 120 to the processing unit 110.

[0046]

Both the application processor 180 and the image processor 190 may include various types of processing devices.

For example, either or both of the application processor 180 and the image processor 190 may include one or more microprocessors, pre-processors (such as image pre-processors), graphics processors, central processing units (CPUs), support circuits, digital signal processors, integrated circuits, memory or any other type of equipment suitable for running applications and suitable for image processing and analysis.

In some embodiments, the application processor 180 and/or the image processor 190 may include any type of single-core or multi-core processors, mobile device microcontrollers, central processing units, and the like.

Various processing devices may be used, including, for example, processors available from manufacturers such as , and may include various architectures (for example, x86 processors, , etc.).

[0047]

In some embodiments, the application processor 180 and/or the image processor 190 may include any EyeQ series processor chips available from .

These processor designs each include multiple processing units with local memory and instruction sets.

Such a processor may include a video input for receiving image data from multiple image sensors, and may also include video output capabilities.

In one embodiment, uses 90 nanometer micron technology running at 332Mhz.

 architecture has two floating-point, hyper-threaded 32-bit RISC CPUs five visual computing engines (VCE), three vector microcode processors Denali 64-bit mobile DDR controller, 128-bit internal supersonic interconnect, dual 16-bit video input and 18-bit video output controller, 16-channel DMA and several peripheral devices.

The MIPS34K CPU manages five VCEs, three VMP.TMs and DMAs, a second MIPS34K CPU and multi-channel DMA, and other peripheral devices.

Five VCEs, three and MIPS34K CPUs can perform the intensive visual calculations required for multi-function package applications.

In another example, can be used in the disclosed embodiment. This is a third-generation processor and is six times more powerful than .

[0048]

Although Figure 1 depicts two separate processing devices included in the processing unit 110, more or fewer processing devices may be used.

For example, in some embodiments, a single processing device may be used to complete the tasks of the application processor 180 and the image processor 190.

In other embodiments, these tasks may be performed by more than two processing devices.

[0049]

The processing unit 110 may include various types of devices.

For example, the processing unit 110 may include various devices, such as a controller, an image preprocessor, a central processing unit (CPU), a support circuit, a digital signal processor, an integrated circuit, a memory, or any other type used for image processing and analysis. device of.

The image preprocessor may include a video processor for capturing, digitizing, and processing images from the image sensor.

The CPU may include any number of microcontrollers or microprocessors.

The support circuit may be any number of circuits generally known in the art, including cache memory, power supply, clock, and input-output circuits.

The memory can store software that controls the operation of the system when executed by the processor.

The memory may include a database and image processing software.

The memory may include any number of random access memory, read only memory, flash memory, disk drive, optical storage, tape storage, removable storage, and other types of storage.

In one example, the memory may be separate from the processing unit 110.

In another example, the memory may be integrated into the processing unit 110.

[0050]

Each memory 140, 150 may include software instructions that, when executed by a processor (for example, the application processor 180 and/or the image processor 190), can control various aspects of the operation of the system 100.

These memory units may include various databases and image processing software.

The memory unit may include random access memory, read only memory, flash memory, disk drive, optical storage, tape storage, removable storage, and/or any other type of storage.

In some embodiments, the memory units 140, 150 may be separated from the application processor 180 and/or the image processor 190.

In other embodiments, these memory units may be integrated into the application processor 180 and/or the image processor 190.

[0051]

The positioning sensor 130 may include any type of device suitable for determining a position associated with at least one component of the system 100.

In some embodiments, the positioning sensor 130 may include a GPS receiver.

Tesla EX1002 Page 1045 IPR2025-00943

Such a receiver can determine the user's position and speed by processing signals propagated by global positioning system satellites. The positioning information from the positioning sensor 130 can be used by the application processor 180 and/or the image processor 190.

[0052]

The user interface 170 may include any device suitable for providing information to one or more users of the system 100 or for receiving input from one or more users of the system 100.

In some embodiments, the user interface 170 may include user input devices, including, for example, a touch screen, a microphone, a keyboard, a pointing device, track wheels, cameras, knobs, buttons, and the like. With such input devices, users can select menu options on the screen by typing instructions or information, providing voice commands, using buttons, pointers, or eye-tracking capabilities, or through any other suitable technology for transferring information to the system 100, Information inputs or commands are provided to the system 100.

[0053]

The user interface 170 may be equipped with one or more processing devices configured to provide information to the user or receive information from the user, and process the information for use by, for example, the application processor 180.

In some embodiments, such a processing device may execute instructions for recognizing and tracking eye movements, receiving and interpreting voice commands, recognizing and interpreting touches and/or gestures made on the touch screen, responding to keyboard input or menu selection, etc. . In some embodiments, the user interface 170 may include a display, a speaker, a haptic device, and/or any other device for providing output information to the user.

[0054]

The map database 160 may include any type of database for storing map data useful to the system 100.

In some embodiments, the map database 160 may include data related to the positioning of various items in a reference coordinate system, including roads, hydrological features, geographic features, businesses, points of interest, restaurants, gas stations, and so on. The map database 160 can store not only the locations of such items, but also descriptors related to these items, including, for example, names associated with any stored features. In some embodiments, the map database 160 may be physically located by other components of the system 100. Alternatively or additionally, the map database 160 or a portion thereof may be remotely located relative to other components of the system 100 (e.g., the processing unit 110). In such an embodiment, the information from the map database 160 may be downloaded through a wired or wireless data connection to the network (for example, through a cellular network and/or the Internet, etc.).

[0055]

The image capture devices 122, 124, and 126 may each include any type of device suitable for capturing at least one image from the environment.

In addition, any number of image capture devices can be used to capture images input to the image processor. Some embodiments may include only a single image capture device, while other embodiments may include two, three, or even four or more image capture devices. The image capturing devices 122, 124, and 126 will be further described below with reference to FIGS. 2B-2E.

[0056]

The system 100 or its various components can be incorporated into a variety of different platforms.

In some embodiments, as shown in FIG. 2A, the system 100 may be included on a vehicle 200. For example, as described above with reference to FIG. 1, the vehicle 200 may be equipped with the processing unit 110 and any other components of the system 100. Although in some embodiments, the vehicle 200 may only be equipped with a single image capture device (eg, a camera), in other embodiments such as those discussed in conjunction with FIGS. 2B-2E, multiple image capture devices may be used. For example, as shown in FIG. 2A, any one of the image capturing devices 122 and 124 of the vehicle 200 may be part of an ADAS (Advanced Driver Assistance System) imaging device.

IPR2025-00943
Tesla EX1002 Page 1046

[0057]

The image capture device included on the vehicle 200 as part of the image capture unit 120 may be located at any suitable location.

In some embodiments, as shown in FIGS. 2A-2E and 3A-3C, the image capture device 122 may be located near the rearview mirror. This location may provide a line of sight similar to that of the driver of the vehicle 200, which may help determine what is visible and invisible to the driver. The image capture device 122 can be positioned anywhere near the rearview mirror, but placing the image capture device 122 on the driver's side of the mirror can further help to obtain an image representing the driver's field of view and/or line of sight.

[0058]

Other positions of the image capturing device of the image capturing unit 120 may also be used.

For example, the image capture device 124 may be located on or in a bumper of the vehicle 200. Such a location may be particularly suitable for image capture devices with a wide field of view. The line of sight of the image capturing device located on the bumper may be different from the line of sight of the driver, and therefore, the bumper image capturing device and the driver may not always see the same object. Image capture devices (e.g., image capture devices 122, 124, and 126) may also be located in other locations. For example, the image capture device may be located on or in one or both of the side mirrors of the vehicle 200, on the roof of the vehicle 200, on the hood of the vehicle 200, on the trunk of the vehicle 200, or on the side of the vehicle 200. , Installed on any window of the vehicle 200, positioned behind or in front of any window of the vehicle 200, and installed in or near the front and/or rear lighting profile of the vehicle 200.

[0059]

In addition to the image capture device, the vehicle 200 may also include various other components of the system 100.

For example, the processing unit 110 may be integrated with the engine control unit (ECU) of the vehicle or separately included on the vehicle 200. The vehicle 200 may also be equipped with a positioning sensor 130 such as a GPS receiver, and may also include a map database 160 and memory units 140 and 150.

[0060]

Figure 2A is a graphical side view representation of an exemplary vehicle imaging system consistent with the disclosed embodiments.

Figure 2B is a graphical top view illustration of the embodiment shown in Figure 2A. As illustrated in FIG. 2B, the disclosed embodiment may include a vehicle 200 that includes a system 100 having a first image capture device 122, a second image capture device 124, and a processing unit 110 in its body. An image capture device 122 is positioned near the rearview mirror and/or close to the driver of the vehicle 200, and the second image capture device 124 is positioned on the bumper area (for example, one of the bumper areas 210) of the vehicle 200 Or among them.

[0061]

As illustrated in FIG. 2C, both the image capture devices 122 and 124 may be positioned near the rearview mirror of the vehicle 200 and/or close to the driver of the vehicle 200.

In addition, although two image capture devices 122 and 124 are shown in FIGS. 2B and 2C, it should be understood that other embodiments may include more than two image capture devices. For example, in the embodiments shown in FIGS. 2D and 2E, the first image capturing device, the second image capturing device, and the third image capturing device 122, 124, and 126 are included in the system 100 of the vehicle 200.

[0062]

As illustrated in FIG. 2D, the image capturing device 122 may be positioned near the rearview mirror of the vehicle 200 and/or close to the driver of the vehicle 200, and the image capturing devices 124 and 126 may be positioned on the bumper of the vehicle 200 Area (e.g., one of the bumper areas 210) or in it.

IPR2025-00943
Tesla EX1002 Page 1047

As shown in FIG. 2E, the image capture devices 122, 124, and 126 may be positioned near the rearview mirror of the vehicle 200 and/or close to the driver of the vehicle 200. The disclosed embodiments are not limited to any specific number of image capture devices and their configurations, and the image capture devices may be positioned in any suitable location within and/or on the vehicle 200.

[0063]

It should be understood that the disclosed embodiments are not limited to vehicles and can be applied in other contexts.

It should also be understood that the disclosed embodiments are not limited to a specific type of vehicle 200, and may be applicable to all types of vehicles, including cars, trucks, trailers, motorcycles, bicycles, self-balancing transportation equipment, and other types of vehicles.

[0064]

The first image capture device 122 may include any suitable type of image capture device.

The image capture device 122 may include an optical axis. In one example, the image capture device 122 may include an Aptina M9V024WVGA sensor with a global shutter. In other embodiments, the image capture device 122 may provide a resolution of 1280x960 pixels and may include a rolling shutter. The image capture device 122 may include various optical elements. In some embodiments, one or more lenses may be included, for example, to provide a desired focal length and field of view for the image capture device. In some embodiments, the image capture device 122 may be associated with a 6mm lens or a 12mm lens. In some embodiments, the image capture device 122 may be configured to capture an image having a desired field of view (FOV) 202 as illustrated in FIG. 2D. For example, the image capture device 122 may be configured to have a regular FOV, such as in the range of 40 degrees to 56 degrees, including 46-degree FOV, 50-degree FOV, 52-degree FOV, or more. Alternatively, the image capturing device 122 may be configured to have a narrow FOV in the range of 23 degrees to 40 degrees, such as a 28 degree FOV or a 36 degree FOV. In addition, the image capturing device 122 may be configured to have a wide FOV in the range of 100 degrees to 180 degrees. In some embodiments, the image capture device 122 may include a wide-angle bumper camera or a wide-angle bumper camera with a FOV of up to 180 degrees.

[0065]

The first image capture device 122 may capture a plurality of first images relative to the scene associated with the vehicle 200.

Each of the plurality of first images may be acquired as a series of image scan lines, and the plurality of first images may be captured using a rolling shutter. Each scan line may include multiple pixels.

[0066]

The first image capture device 122 may have a scan rate associated with the acquisition of each of the first series of image scan lines.

The scan rate may refer to the rate at which the image sensor can collect image data associated with each pixel included in a specific scan line.

[0067]

The image capture devices 122, 124, and 126 may include any suitable type and number of image sensors, including, for example, CCD sensors or CMOS sensors.

In one embodiment, a CMOS image sensor can be used in conjunction with a rolling shutter, so that each pixel in a row is read one at a time, and the scanning of the rows is performed on a row-by-row basis until the entire image frame is captured. In some embodiments, the rows may be captured sequentially from top to bottom with respect to the frame.

[0068]

The use of a rolling shutter may cause different rows of pixels to be exposed and captured at different times, which may cause skew and other image artifacts in the captured image frame.

On the other hand, when the image capture device 122 is configured to operate with a global shutter or a synchronized shutter, all pixels can be exposed within the same amount of time and during a common exposure period. Therefore, the image data in the frames collected from the system employing the global shutter represents a snapshot of the entire FOV (such as FOV 202) at a specific time. In contrast, in rolling shutter applications, each row in the frame is exposed at a different time, and the data is captured at a different time. Therefore, the moving object may be distorted in an image capturing device with a rolling shutter. This phenomenon will be described in more detail below.

[0069]

The second image capturing device 124 and the third image capturing device 126 may be any types of image capturing devices.

Similar to the first image capture device 122, each of the image capture devices 124 and 126 may include an optical axis. In one example, each of the image capture devices 124 and 126 may include an Aptina M9V024WVGA sensor with a global shutter. Alternatively, each of the image capture devices 124 and 126 may include a rolling shutter. Similar to the image capture device 122, the image capture devices 124 and 126 may be configured to include various lenses and optical elements. In some embodiments, the lenses associated with image capture devices 124 and 126 may provide FOVs (such as FOVs 204 and 206) that are the same or different (e.g., FOVs) associated with image capture device 122 (such as FOV 202) ,Narrower). For example, the image capture devices 124 and 126 may have a FOV of 40 degrees, 30 degrees, 26 degrees, 23 degrees, 20 degrees, or lower.

[0070]

The image capture devices 124 and 126 may capture a plurality of second and third images relative to the scene associated with the vehicle 200.

Each of the plurality of second images and third images may be collected as the second series of image scan lines and the third series of image scan lines, and each of the plurality of second images and third images may use a rolling shutter To capture. Each scan line or scan line may have multiple pixels. The image capturing devices 124 and 126 may have a second scan rate and a third scan rate associated with the acquisition of each image scan line included in the second series and the third series.

[0071]

Each image capture device 122, 124, and 126 may be positioned at any suitable position and orientation relative to the vehicle 200.

The relative orientation of the image capture devices 122, 124, and 126 can be selected to help fuse together the information collected from the image capture devices. For example, in some embodiments, the FOV associated with the image capture device 124 (such as FOV 204) may be the FOV associated with the image capture device 122 (such as FOV 202) and the FOV associated with the image capture device 126 (such as FOV 206).) Overlap partially or completely.

[0072]

The image capture devices 122, 124, and 126 may be located at any suitable relative height on the vehicle 200.

In one example, there is a height difference between the image capture devices 122, 124, and 126, which can provide sufficient disparity information to enable stereo analysis. For example, as shown in Figure 2A, the two image capture devices 122 and 124 are at different heights. There may also be a lateral displacement difference between the image capturing devices 122, 124, and 126, for example, additional disparity information for stereo analysis is given by the processing unit 110. As shown in Figure 2C and Figure 2D, the difference in lateral displacement can be expressed as dx. In some embodiments, the front displacement and the rear displacement (eg, range displacement) may exist between the image capture devices 122, 124, and 126. For example, the image capture device 122 may be located 0.5 to 2 meters or more behind the image capture device 124 and/or the image capture device 126. This type of displacement enables one of the image capture devices to cover the potential blind spot of the other image capture device.

[0073]

IPR2025-00943
Tesla EX1002 Page 1049

The image capture device 122 may have any suitable resolution capability (for example, the number of pixels associated with the image sensor), and the resolution of the image sensor associated with the image capture device 122 may be higher or lower than that of the image capture device. The resolution of the image sensor associated with 124 and 126, or the same.

In some embodiments, the image sensors associated with image capture device 122 and/or image capture devices 124 and 126 may have a resolution of 640x480, 1024x768, 1280x960, or any other suitable resolution.

[0074]

The frame rate (eg, the rate at which the image capture device collects a set of pixel data for one image frame before moving to capture the pixel data associated with the next image frame) may be controllable.

The frame rate associated with image capture device 122 may be higher, lower than, or the same as the frame rate associated with image capture devices 124 and 126. The frame rate associated with image capture devices 122, 124, and 126 may depend on various factors that may affect the timing of the frame rate. For example, one or more of the image capture devices 122, 124, and 126 may include selectable pixel delay periods that are set in one or more of the image sensors in the acquisition and image capture devices 122, 124, and/or 126. The image data associated with each pixel is applied before or after. Generally, the image data corresponding to each pixel can be collected according to the clock rate of the device (for example, one pixel per clock cycle). In addition, in embodiments that include rolling shutters, one or more of the image capture devices 122, 124, and 126 may include selectable horizontal blanking periods, which may be selected between the acquisition and image capture devices 122, 124, and 126. The image data associated with the pixel rows of the image sensor in 124 and/or 126 is applied before or after. In addition, one or more of the image capture devices 122, 124, and/or 126 may include a selectable vertical blanking period, which may be used to capture and capture image frames of the image capture devices 122, 124, and 126. The associated image data is applied before or after.

[0075]

These timing controls can synchronize the frame rates associated with the image capture devices 122, 124, and 126 even when the line scan rate of each is different.

In addition, as will be discussed in more detail below, in addition to other factors (e.g., image sensor resolution, maximum line scan rate, etc.), these selectable timing controls can also make the FOV from the image capture device 122 and the image capture device 122 The image capture in the area where one or more FOVs of 124 and 126 overlap can be synchronized even when the field of view of the image capturing device 122 is different from the FOV of the image capturing devices 124 and 126.

[0076]

The frame rate timing in the image capture devices 122, 124, and 126 may depend on the resolution of the associated image sensor.

For example, assuming that two devices have similar line scan rates, if one device includes an image sensor with a resolution of 640x480, and the other device includes an image sensor with a resolution of 1280x960, then capture from the sensor with a higher resolution The frame of image data will require more time.

[0077]

Another factor that can affect the timing of image data acquisition in image capture devices 122, 124, and 126 is the maximum line scan rate.

For example, the collection of lines of image data from the image sensors included in the image capture devices 122, 124, and 126 will require some minimum amount of time. Assuming that the pixel delay period is not increased, this minimum amount of time for the lines of image data to be collected will be related to the maximum line scan rate of the particular device. A device that is given a higher maximum line scan rate may provide a higher frame rate than a device that has a lower maximum line scan rate. In some embodiments, one or more of the image capture devices 124 and 126 may have a maximum line scan rate that is higher than the maximum line scan rate associated with the image capture device 122. In some embodiments, the maximum line scan rate of the image capture device 124 and/or 126 may be 1.25 times, 1.5 times, 1.75 times, or 2 times or more times the maximum line scan rate of the image capture device 122.

[0078]

IPR2025-00943
Tesla EX1002 Page 1050

In another embodiment, the image capture devices 122, 124, and 126 may have the same maximum line scan rate, but the image capture device 122 may operate at a scan rate less than or equal to its maximum scan rate.

The system can be configured such that one or more image capture devices 124 and 126 operate at a line scan rate equal to the line scan rate of image capture device 122. In other examples, the system may be configured such that the line scan rate of the image capture device 124 and/or the image capture device 126 may be 1.25, 1.5, 1.75, or 2 times the line scan rate of the image capture device 122, or More times.

[0079]

In some embodiments, the image capture devices 122, 124, and 126 may be asymmetrical.

That is, they can include cameras with different fields of view (FOV) and focal lengths. For example, the field of view of the image capture devices 122, 124, and 126 may include any desired area relative to the environment of the vehicle 200. In some embodiments, one or more of the image capture devices 122, 124, and 126 may be configured to collect image data from the environment in front of the vehicle 200, behind the vehicle 200, on the side of the vehicle 200, or a combination thereof.

[0080]

In addition, the focal length associated with each image capture device 122, 124, and/or 126 may be selectable (for example, by including an appropriate lens, etc.) so that each device captures at a desired distance range relative to the vehicle 200 Image of the object.

For example, in some embodiments, the image capture devices 122, 124, and 126 may capture images of close objects within a few meters of the vehicle. The image capture devices 122, 124, and 126 may also be configured to capture images of objects at a distance from the vehicle (e.g., 25m, 50m, 100m, 150m or more). In addition, the focal lengths of the image capturing devices 122, 124, and 126 can be selected so that one image capturing device (for example, the image capturing device 122) can capture images of objects relatively close to the vehicle (for example, within 10 m or within 20 m), and Other image capture devices (e.g., image capture devices 124 and 126) may capture images of objects farther from the vehicle 200 (e.g., greater than 20m, 50m, 100m, 150m, etc.).

[0081]

According to some embodiments, the FOV of one or more image capture devices 122, 124, and 126 may have a wide angle.

For example, having a FOV of 140 degrees may be advantageous, particularly for image capture devices 122, 124, and 126 that can be used to capture images of the area near the vehicle 200. For example, the image capture device 122 may be used to capture an image of an area on the right or left side of the vehicle 200, and in such an embodiment, the image capture device 122 may be required to have a wide FOV (e.g., at least 140 degrees).

[0082]

In some embodiments, the system may define the "nearby" of the vehicle based on predetermined restrictions or dynamically determined restrictions.

The vicinity of the vehicle, for example, can be considered to be within a certain distance of the vehicle in any one or more directions (such as 1 meter, 5 meters, 10 meters, 25 meters, 50 meters, 100 meters, 250 meters, 500 meters, or 1000 meters). Any area. In some embodiments, the area defined as the vicinity of the vehicle may be irregularly shaped so that it extends different distances in different directions; the vicinity in the forward direction defined by the driving of the vehicle may extend, for example, several hundred meters , And the vicinity in the rearward or lateral direction can only extend for tens of meters or less, for example. In some embodiments, the proximity of the defined vehicle may dynamically change based on information captured by one or more image capture devices, so that the proximity may change, for example, based on which parts of the road are visible in the FOV of the image capture device. In some embodiments, the vicinity of the vehicle may dynamically change, for example, according to any perceivable conditions or states of the road or the vehicle, including vehicle speed, posted rate limits, road conditions, road congestion, and road size. , Visibility distance, ambient brightness, time, the presence of other vehicles, the presence of pedestrians or any other suitable factors.

IPR2025-00943
Tesla EX1002 Page 1051

[0083]

The field of view associated with each of the image capture devices 122, 124, and 126 may depend on the respective focal length.

For example, as the focal length increases, the corresponding field of view decreases.

[0084]

The image capture devices 122, 124, and 126 may be configured to have any suitable field of view.

In a specific example, the image capture device 122 may have a horizontal FOV of 46 degrees, the image capture device 124 may have a horizontal FOV of 23 degrees, and the image capture device 126 may have a horizontal FOV between 23 degrees and 46 degrees. In another example, the image capture device 122 may have a horizontal FOV of 52 degrees, the image capture device 124 may have a horizontal FOV of 26 degrees, and the image capture device 126 may have a horizontal FOV between 26 degrees and 52 degrees. In some embodiments, the ratio of the FOV of the image capture device 122 to the FOV of the image capture device 124 and/or the image capture device 126 may vary from 1.5 to 2.0. In other embodiments, the ratio can vary between 1.25 and 2.25.

[0085]

The system 100 may be configured such that the field of view of the image capture device 122 and the field of view of the image capture device 124 and/or the image capture device 126 at least partially or fully overlap.

In some embodiments, the system 100 may be configured such that the field of view of the image capture devices 124 and 126 is, for example, within (e.g., narrower) of the field of view of the image capture device 122 and shares a common center therewith. In other embodiments, the image capture devices 122, 124, and 126 may capture adjacent FOVs, or may have partial overlaps in their FOVs. In some embodiments, the field of view of the image capture devices 122, 124, and 126 can be aligned so that the center of the narrower FOV image capture device 124 and/or 126 can be located in the lower half of the field of view of the wider FOV device 122 part.

[0086]

Figure 2F is a graphical representation of an exemplary vehicle control system consistent with the disclosed embodiments.

As shown in FIG. 2F, the vehicle 200 may include a throttle system 220, a braking system 230, and a steering system 240. The system 100 can be connected to one or more of the throttle system 220, the braking system 230, and the steering system 240 through one or more data links (for example, any wired and/or wireless links or links used to transmit data). Multiple provide inputs (e.g., control signals). For example, based on the analysis of the images captured by the image capture devices 122, 124, and/or 126, the system 100 may provide control signals to one or more of the throttle system 220, the braking system 230, and the steering system 240 to navigate the vehicle 200 (for example, by causing acceleration, turning, lane deviation, etc.). In addition, the system 100 may receive input from one or more of the flow system 220, the braking system 230, and the steering system 240 that indicate the operating conditions of the vehicle 200 (eg, speed, whether the vehicle 200 is braking and/or turning, etc.) . Further details are provided below in conjunction with Figures 4-7.

[0087]

As shown in FIG. 3A, the vehicle 200 may also include a user interface 170 for interacting with the driver or passengers of the vehicle 200.

For example, the user interface 170 in a vehicle application may include a touch screen 320, a knob 330, a button 340, and a microphone 350. The driver or passenger of the vehicle 200 can also use handles (for example, on or near the steering column of the vehicle 200, including turn signal handles, for example), buttons (for example, on the steering wheel of the vehicle 200), etc., to interact with the system 100 interactions. In some embodiments, the microphone 350 may be positioned adjacent to the rear view mirror 310. Similarly, in some embodiments, the image capture device 122 may be located near the rear view mirror 310. In some embodiments, the user interface 170 may also include one or more speakers 360 (eg, speakers of a vehicle audio system). For example, the system 100 may provide various notifications (e.g., alarms) via the speaker 360.

[0088]

3B-3D are illustrations of an exemplary camera base 370 configured to be positioned behind a rearview mirror (eg, rearview mirror 310) and facing the windshield of a vehicle, consistent with the disclosed embodiments.

As shown in FIG. 3B, the camera dock 370 may include image capture devices 122, 124, and 126. The image capture devices 124 and 126 may be positioned behind the anti-glare plate 380, which may be flush with the vehicle windshield and include a combination of films and/or anti-reflective materials. For example, the anti-glare plate 380 may be positioned so that it is aligned against a vehicle windshield with a matching slope. In some embodiments, each of the image capture devices 122, 124, and 126 may be positioned behind the anti-glare plate 380, for example as depicted in FIG. 3D. The disclosed embodiments are not limited to any specific configuration of the image capture devices 122, 124, and 126, the camera base 370, and the anti-glare plate 380. FIG. 3C is an illustration of the camera base 370 shown in FIG. 3B from a front perspective.

[0089]

Those skilled in the art with the benefits of the present disclosure will recognize that various changes and/or modifications can be made to the above-disclosed embodiments.

For example, not all components are necessary for the operation of the system 100. In addition, any component may be located in any suitable part of the system 100, and the components may be rearranged into various configurations while providing the functions of the disclosed embodiments. Therefore, the above configuration is an example, and regardless of the configuration discussed above, the system 100 can provide a wide range of functions to analyze the surrounding environment of the vehicle 200, and in response to the analysis, navigate and/or otherwise control and/or operate the vehicle 200. The navigation, control, and/or operation of the vehicle 200 may include (directly or via an intermediate controller such as the above-mentioned controller) enabling and/or disabling various features, components, devices, modes, System and/or subsystem. Alternatively or additionally, navigation, control, and/or operation may include, for example, by providing visual, audio, tactile, and/or other sensory alerts and/or instructions to users, drivers, and passengers that may be located inside or outside the vehicle 200. , Passersby and/or other vehicles or user interactions.

[0090]

As discussed in further detail below and consistent with various disclosed embodiments, the system 100 may provide various features related to autonomous driving and/or driver assistance technology.

For example, the system 100 may analyze image data, positioning data (for example, GPS location information), map data, speed data, and/or data from sensors included in the vehicle 200. The system 100 may collect data for analysis, for example, from the image acquisition unit 120, the positioning sensor 130, and other sensors. In addition, the system 100 can analyze the collected data to determine whether the vehicle 200 should take a certain action, and then automatically take the determined action without human intervention. For example, when the vehicle 200 is navigating without human intervention, (for example, by sending a control signal to one or more of the throttle system 220, the braking system 230, and the steering system 240), the system 100 can automatically control the vehicle 200 braking, acceleration and/or steering. In addition, the system 100 can analyze the collected data, and based on the analysis of the collected data, issue warnings, instructions, suggestions, alarms, or instructions to the driver, passengers, users, or other people inside or outside the vehicle (or other vehicles). .

[0091]

In addition, consistent with the disclosed embodiments, the functions provided by the system 100 enable the vehicle 200 to take different actions to navigate the vehicle 200 in the lane and/or relative to other vehicles and/or objects.

For example, the system 100 can adjust the positioning of the vehicle 200 relative to the lane in which the vehicle 200 is traveling and/or relative to objects located near the vehicle 200, select a specific lane that the vehicle 200 uses while traveling, and respond to intrusive A vehicle, such as a vehicle that is trying to move into the lane in which the vehicle 200 is traveling, takes action. In addition, the system 100 can control the speed of the vehicle 200 in different scenarios, such as when the vehicle 200 is turning. The system 100 allows the vehicle 200 to imitate the actions of the leading vehicle or monitor the target vehicle and navigate the vehicle 200 so that the vehicle passes the target vehicle. Additional details related to the various embodiments provided by the system 100 are provided below.

[0092]

IPR2025-00943
Tesla EX1002 Page 1053

Front-viewing multi-imaging system

[0093]

As discussed above, the system 100 may provide a driving assistance function using a multi-camera system.

The multi-camera system may use one or more cameras facing the forward direction of the vehicle.

In other embodiments, the multi-camera system may include one or more cameras facing the side of the vehicle or the rear of the vehicle. In one embodiment, for example, the system 100 may use a two-camera imaging system, where a first camera and a second camera (e.g., image capture devices 122 and 124) may be positioned in front of a vehicle (e.g., vehicle 200) and/ Or side. The first camera may have a field of view larger than, smaller than, or partially overlapped with the field of view of the second camera. In addition, the first camera may be connected to the first image processor to perform monocular image analysis on the images provided by the first camera, and the second camera may be connected to the second image processor to perform monocular image analysis on the images provided by the second camera Monocular image analysis. The outputs (for example, processed information) of the first image processor and the second image processor may be combined. In some embodiments, the second image processor may receive images from both the first camera and the second camera to perform stereo analysis. In another embodiment, the system 100 may use a three-camera imaging system in which each of the cameras has a different field of view. Therefore, such a system can make decisions based on information obtained from objects located at varying distances in both the front and side of the vehicle. Reference to monocular image analysis may refer to an example of performing image analysis based on an image captured from a single viewpoint (for example, from a single camera). Stereoscopic image analysis may refer to an example in which image analysis is performed based on two or more images captured through one or more changes in image capturing parameters. For example, captured images suitable for performing stereoscopic image analysis may include images captured from two or more different positions, images captured from different fields of view using different focal lengths together with parallax information, and the like.

[0094]

For example, in one embodiment, the system 100 may implement a three-camera configuration using image capture devices 122-126.

In such a configuration, the image capture device 122 can provide a narrow field of view (for example, 34 degrees or other values selected from the range of about 20 degrees to 45 degrees, etc.), and the image capture device 124 can provide a wide field of view (for example, 150 Degrees or other values selected from the range of about 100 degrees to about 180 degrees), and the image capture device 126 may provide an intermediate field of view (eg, 46 degrees or other values selected from the range of about 35 degrees to about 60 degrees). In some embodiments, the image capture device 126 may be used as the main camera or the primary camera. The image capture devices 122-126 may be positioned behind the rear view mirror 310, and may be positioned substantially side by side (e.g., 6 cm apart). Furthermore, in some embodiments, as discussed above, one or more image capture devices 122-126 may be installed behind the anti-glare plate 380 that is flush with the windshield of the vehicle 200. Such protection can be used to minimize the impact of any reflections from the interior of the car on the image capture devices 122-126.

[0095]

In another embodiment, as discussed above in connection with FIGS. 3B and 3C, a wide field of view camera (for example, the image capture device 124 in the above example) may be installed below the narrow field of view camera and the main field of view camera (For example, the image capture devices 122 and 126 in the above example).

This configuration can provide a line of sight that is not affected by a wide-field camera. In order to reduce reflections, the camera may be installed close to the windshield of the vehicle 200, and a polarizer may be included on the camera to attenuate the reflected light.

[0096]

The three-camera system can provide certain performance characteristics.

IPR2025-00943
Tesla EX1002 Page 1054

IPR2025-00943 Tesla EX1002 Page 1055

For example, some embodiments may include the ability to verify the detection of an object by one camera based on the detection result from another camera. In the three-camera configuration discussed above, the processing unit 110 may, for example, include three processing devices (for example, three EyeQ series processor chips as discussed above), each of which is dedicated to processing by one or more images captured by image capture devices 122-126.

[0097]

In the three-camera system, the first processing device can receive images from both the main camera and the narrow field of view camera, and perform the visual processing of the narrow FOV camera, for example, to detect other vehicles, pedestrians, lane markings, traffic signs, and traffic Lamps and other road objects.

In addition, the first processing device may calculate the parallax of pixels between the image from the main camera and the image from the narrow camera to create a 3D reconstruction of the environment of the vehicle 200. Subsequently, the first processing device may combine the 3D reconstruction with 3D map data (for example, a depth map) or with 3D information calculated based on information from another camera.

[0098]

The second processing device can receive images from the main camera and perform visual processing to detect other vehicles, pedestrians, lane markings, traffic signs, traffic lights, and other road objects.

In addition, the second processing device may calculate the camera displacement, and based on the displacement, calculate the disparity of pixels between consecutive images, and create a 3D reconstruction of the scene (for example, a motion recovery structure). The second processing device may send the 3D reconstruction based on the motion restoration structure to the first processing device to be combined with the stereoscopic 3D image or with the depth information obtained through stereo processing.

[0099]

The third processing device may receive the image from the wide FOV camera and process the image to detect vehicles, pedestrians, lane markings, traffic signs, traffic lights, and other road objects.

The third processing device can execute additional processing instructions to analyze the image to identify objects moving in the image, such as vehicles that change lanes, pedestrians, and the like.

[0100]

In some embodiments, image stream-based information that is separately captured and processed may provide opportunities for providing redundancy in the system.

Such redundancy may include, for example, using a first image capturing device and an image processed by the device to verify and/or supplement information obtained by capturing and processing image information from at least a second image capturing device.

[0101]

In some embodiments, the system 100 may use two image capture devices (e.g., image capture devices 122 and 124) to provide navigation assistance to the vehicle 200, and may use a third image capture device (e.g., image capture device 126) to provide navigation assistance to the vehicle 200. Provide redundancy and verify the analysis of the data received from the other two image capture devices.

For example, in such a configuration, the image capture devices 122 and 124 may provide images for stereo analysis by the system 100 to navigate the vehicle 200, and the image capture device 126 may provide images for monocular analysis by the system 100, To provide redundancy and verification based on information obtained from the image captured by the image capture device 122 and/or the image capture device 124. That is, the image capture device 126 (and the corresponding processing device) can be considered to provide a redundant subsystem for checking the analysis obtained from the image capture devices 122 and 124 (for example, to provide automatic emergency braking (AEB) system).

[0102]

Those skilled in the art will recognize that the above camera configuration, camera placement, number of cameras, camera positions, etc. are just examples.

These components and others described with respect to the overall system can be assembled and used in a variety of different configurations without departing from the scope of the disclosed embodiments. Further details concerning the use of the multi-camera system for providing driver assistance and/or autonomous vehicle functions are as follows.

[0103]

4 is an exemplary functional block diagram of the memory 140 and/or 150, which can store/program instructions for executing one or more operations consistent with the disclosed embodiments.

Although the memory 140 is referred to below, those skilled in the art will recognize that instructions may be stored in the memory 140 and/or 150.

[0104]

As shown in FIG. 4, the memory 140 may store a monocular image analysis module 402, a stereo image analysis module 404, a speed and acceleration module 406, and a navigation response module 408.

The disclosed embodiment is not limited to any specific configuration of the memory 140. In addition, the application processor 180 and/or the image processor 190 may execute instructions stored in any one of the modules 402-408 included in the memory 140. Those skilled in the art will understand that the reference to the processing unit 110 in the following discussion may refer to the application processor 180 and the image processor 190 individually or collectively. Therefore, any step in the following process can be executed by one or more processing devices.

[0105]

In one embodiment, the monocular image analysis module 402 may store instructions (such as computer vision software), which, when executed by the processing unit 110, will The image is subjected to monocular image analysis.

In some embodiments, the processing unit 110 may combine information from a set of images with additional sensory information (for example, information from radar) to perform monocular image analysis. As described below in conjunction with FIGS. 5A-5D, the monocular image analysis module 402 may include a set of features used to detect the set of images (such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic Lights, dangerous objects, and any other features associated with the vehicle's environment). In some embodiments, for any given frame, the environment of the vehicle may include an area or space (the FOV of the imaging device) covered by any one of the imaging devices or by any combination thereof. As discussed below in connection with the navigation response module 408, based on this analysis, the system 100 (eg, via the processing unit 110) can cause one or more navigation responses in the vehicle 200, such as turning, lane shifting, acceleration changes, and so on. In addition, based on this analysis, the system 100 (eg, via the processing unit 110) may cause one or more navigation and/or vehicle control responses, including but not limited to controlling the vehicle 200 and any associated systems and/or systems discussed above Various ways of subsystems. In some embodiments, any or all of the above-mentioned responses or actions may be referred to as navigation actions or navigation responses for the sake of simplicity.

[0106]

In one embodiment, the stereoscopic image analysis module 404 may store instructions (such as computer vision software) that, when executed by the processing unit 110, capture images from any one of the image capture devices 122, 124, and 126. The combination of the first set of images and the second set of images collected by the device performs stereoscopic image analysis.

In some embodiments, the processing unit 110 may combine information from the first set of images and the second set of images with additional sensory information (for example, information from radar) to perform stereoscopic image analysis. For example, the stereoscopic image analysis module 404 may include instructions for performing stereoscopic image analysis based on the first set of images captured by the image capture device 124 and the second set of images captured by the image capture device 126. As described below in conjunction with FIG. 6, the stereo image analysis module 404 may include a set of

IPR2025-00943
Tesla EX1002 Page 1056

features (such as lane markings, vehicles, pedestrians, road signs, highway exits) in the first set of images and the second set of images. Ramps, traffic lights, dangerous objects, etc.) instructions. As discussed below in conjunction with the navigation response module 408, based on this analysis, the processing unit 110 may cause one or more navigation responses in the vehicle 200, such as turning, lane shifting, changes in acceleration, and the like.

[0107]

In one embodiment, the speed and acceleration module 406 may store data configured to receive data from one or more computing devices and electromechanical devices in the vehicle 200 that are configured to cause a change in the speed and/or acceleration of the vehicle 200 Perform analysis.

For example, the processing unit 110 may execute instructions associated with the speed and acceleration module 406 to calculate the target speed of the vehicle 200 based on the data obtained by the execution of the monocular image analysis module 402 and/or the stereo image analysis module 404. Such data may include, for example, target location, speed and/or acceleration, location and/or velocity of vehicle 200 relative to nearby vehicles, pedestrians or road objects, and location information of vehicle 200 relative to the lane markings of the road. In addition, the processing unit 110 may calculate the vehicle 200 based on sensory input (for example, information from radar) and input from other systems of the vehicle 200 (such as the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200). The target rate. Based on the calculated target rate, the processing unit 110 may transmit electrical signals to the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200 to be triggered, for example, by physically depressing the brake or releasing the accelerator of the vehicle 200 Changes in speed and/or acceleration.

[0108]

In one embodiment, the navigation response module 408 may store software executable by the processing unit 110 to determine the required navigation response based on the data obtained by the execution of the monocular image analysis module 402 and/or the stereo image analysis module 404.

Such data may include location information and speed information associated with nearby vehicles, pedestrians, and road objects, target location information of the vehicle 200, and the like. In addition, in some embodiments, the navigation response may be (partially or fully) based on map data, the predetermined location of the vehicle 200, and/or the relationship between the vehicle 200 and the execution of the monocular image analysis module 402 and/or the stereo image analysis module 404. The relative velocity or relative acceleration between the detected one or more objects. The navigation response module 408 may also determine the required navigation based on sensory input (e.g., information from radar) and input from other systems of the vehicle 200 (such as the throttle system 220, braking system 230, and steering system 240 of the vehicle 200). response. Based on the required navigation response, the processing unit 110 may transmit electrical signals to the throttle system 220, braking system 230, and steering system 240 of the vehicle 200 to, for example, turn the steering wheel of the vehicle 200 to achieve a predetermined angle of rotation to trigger the desired Navigation response. In some embodiments, the processing unit 110 may use the output of the navigation response module 408 (eg, the required navigation response) as an input to the execution of the speed and acceleration module 406 to calculate the change in the speed of the vehicle 200.

[0109]

FIG. 5A is a flowchart showing an exemplary process 500A for eliciting one or more navigation responses based on monocular image analysis consistent with the disclosed embodiments.

In step 510, the processing unit 110 may receive a plurality of images via the data interface 128 between the processing unit 110 and the image acquisition unit 120. For example, a camera included in the image capture unit 120 (such as an image capture device 122 having a field of view 202) can capture multiple images of the forward area of the vehicle 200 (or, for example, the side or the rear of the vehicle), and pass The data connection (for example, digital, wired, USB, wireless, Bluetooth, etc.) transmits them to the processing unit 110. As described in further detail below in conjunction with FIGS. 5B-5D, in step 520, the processing unit 110 may cause the monocular image analysis module 402 to perform analysis on multiple images. By performing this analysis, the processing unit 110 can detect a set of features in the set of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, and so on.

[0110]

In step 520, the processing unit 110 may also enable the monocular image analysis module 402 to perform detection of various road dangerous objects (parts of truck tires, fallen road signs, loose goods, small animals, etc.).

Road hazards may vary in structure, shape, size, and color, which may make decisions about such hazards more challenging. In some embodiments, the processing unit 110 may enable the monocular image analysis module 402 to perform multi-frame analysis on multiple images to detect road hazards. For example, the processing unit 110 may estimate the camera motion between consecutive image frames, and calculate the disparity of pixels between the frames to construct a road or certain areas of the road (for example, a suspected road hazard defined around the image data). 3D map of the object within the area of interest). Subsequently, the processing unit 110 may use the 3D map to detect road hazards. For example, segmentation technology and classification technology can be used to detect road surfaces and hazards that exist on the road surface. For example, as explained in additional detail below, the system can detect features of the road surface, such as drawn lane lines or lane markings, and additionally can detect features of roads that exist above the road surface, such as cones/cylinders in the construction area Body/poles/obstacles, guardrails, other vehicles, road signs, etc.

[0111]

In step 530, the processing unit 110 may cause the navigation response module 408 to perform one or more navigation responses in the vehicle 200 based on the analysis performed in step 520 and the technique described above in conjunction with FIG. 4.

For example, the navigation response may include turns, lane shifts, changes in acceleration, and so on. In some embodiments, the processing unit 110 may use data derived from the execution of the speed and acceleration module 406 to cause one or more navigation responses. In addition, multiple navigation responses can occur simultaneously, sequentially, or any combination of them. For example, the processing unit 110 may cause the vehicle 200 to move across a lane, and then accelerate, for example, by sequentially transmitting control signals to the steering system 240 and the throttle system 220 of the vehicle 200. Alternatively, the processing unit 110 may brake the vehicle 200 while shifting the lane, for example, by simultaneously transmitting control signals to the braking system 230 and the steering system 240 of the vehicle 200.

[0112]

Figure 5B is a flowchart showing an exemplary process 500B for detecting one or more vehicles and/or pedestrians in a set of images consistent with the disclosed embodiments.

The processing unit 110 may cause the monocular image analysis module 402 to execute the process 500B. In step 540, the processing unit 110 may determine a set of candidate objects representing possible vehicles and/or pedestrians. For example, the processing unit 110 may scan one or more images, compare the images with one or more predetermined patterns, and identify possible locations that may contain objects of interest (for example, vehicles, pedestrians, or parts of them) Every image. The predetermined pattern can be designed in such a way to achieve a high rate of "false hits" and a low rate of "missing". For example, the processing unit 110 may use a low threshold similar to a predetermined pattern for identifying vehicles or pedestrians that are possible candidates. Doing so may enable the processing unit 110 to reduce the possibility of missing (for example, failing to recognize) candidates representing vehicles or pedestrians.

[0113]

In step 542, the processing unit 110 may filter the set of candidate objects based on the classification criteria to exclude certain candidates (for example, irrelevant or less relevant objects).

Such criteria can be derived from various attributes associated with object types stored in a database (e.g., a database stored in memory 140). The attributes may include the shape, size, texture, location (for example, relative to the vehicle 200) of the object, and the like. Therefore, the processing unit 110 may use one or more sets of criteria to reject erroneous candidates in the set of candidate objects.

[0114]

In step 544, the processing unit 110 may analyze multiple frames of images to determine whether an object in the set of candidate objects represents a vehicle and/or a pedestrian.

For example, the processing unit 110 may track candidate objects detected in consecutive frames, and accumulate frame-by-frame data associated with the detected objects (for example, the size, positioning, etc. relative to the vehicle 200). In addition, the processing unit 110 may estimate the parameters of the detected object, and compare the frame-by-frame positioning data of the object with the predicted positioning.

[0115]

In step 546, the processing unit 110 may construct a set of measurement values about the detected object.

Such measurement values may include, for example, the size, positioning, speed, and acceleration values (relative to the vehicle 200) associated with the detected object. In some embodiments, the processing unit 110 may be based on estimation techniques that use a series of time-based observations (such as a Kalman filter or linear quadratic estimation (LQE)), and/or based on different object types (e.g., cars, Modeling data of trucks, pedestrians, bicycles, road signs, etc.) can be used to construct measurement values. The Kalman filter may be based on a measurement of the ratio of the object, where the measurement of the ratio is proportional to the collision time (e.g., the amount of time the vehicle 200 reaches the object). Therefore, by performing steps 540-546, the processing unit 110 can recognize vehicles and pedestrians appearing in the group of captured images, and derive information associated with the vehicles and pedestrians (for example, location, speed, size). As described above in conjunction with FIG. 5A, based on the recognized and derived information, the processing unit 110 may cause one or more navigation responses in the vehicle 200.

[0116]

In step 548, the processing unit 110 may perform optical flow analysis of one or more images to reduce the possibility of detecting "false hits" and missing candidate objects representing vehicles or pedestrians.

For example, optical flow analysis may refer to analyzing the movement pattern of the vehicle 200 in one or more images associated with other vehicles and pedestrians, and the movement pattern is different from the movement of the road surface. The processing unit 110 may calculate the motion of the candidate object by observing different positions of the object in a plurality of image frames captured at different times. The processing unit 110 may use the values of positioning and time as inputs for calculating a mathematical model of the motion of the candidate object. Therefore, optical flow analysis may provide another method for detecting vehicles and pedestrians near the vehicle 200. The processing unit 110 may perform optical flow analysis in combination with steps 540-546 to provide redundancy for detecting vehicles and pedestrians and improve the reliability of the system 100.

[0117]

Figure 5C is a flowchart showing an exemplary process 500C for detecting road markings and/or lane geometry information in a set of images consistent with the disclosed embodiments.

The processing unit 110 can execute the monocular image analysis module 402 to implement the process 500C. In step 550, the processing unit 110 may detect a group of objects by scanning one or more images. In order to detect lane markings, lane geometry information, and other related road marking segments, the processing unit 110 may filter the group of objects to exclude those objects that are determined to be irrelevant (for example, small pits, small rocks, etc.). In step 552, the processing unit 110 may combine the segments that belong to the same road marking or lane marking detected in step 550. Based on this grouping, the processing unit 110 may develop a model representing the detected segment, such as a mathematical model.

[0118]

In step 554, the processing unit 110 may construct a set of measurements associated with the detected segment.

In some embodiments, the processing unit 110 may create a projection of the segment detected from the image plane onto the actual plane. The projection can be characterized using a cubic polynomial with coefficients corresponding to physical properties such as the location, slope, curvature, and curvature derivative of the detected road. When generating the projection, the processing unit 110 may consider the change of the road surface and the rate of pitch and roll associated with the vehicle 200. In addition, the processing unit 110 may model the road elevation by analyzing existing positioning and movement clues on the road surface. In addition, the processing unit 110 may estimate the rate of pitch and roll associated with the vehicle 200 by tracking a set of feature points in one or more images.

IPR2025-00943
Tesla EX1002 Page 1059

[0119]

In step 556, the processing unit 110 may perform multi-frame analysis, for example, by tracking the detected segments in consecutive image frames and accumulating frame-by-frame data associated with the detected segments.

As the processing unit 110 performs multi-frame analysis, the set of measurement values constructed in step 554 may become more reliable and associated with higher and higher confidence levels. Therefore, by performing steps 550-556, the processing unit 110 can recognize the road markings appearing in the group of captured images, and obtain information about the lane geometry. As described above in conjunction with FIG. 5A, based on the recognized and derived information, the processing unit 110 may cause one or more navigation responses in the vehicle 200.

[0120]

In step 558, the processing unit 110 may consider additional information sources to further develop a safety model for the vehicle 200 in the context of the surrounding environment of the vehicle 200.

The processing unit 110 may use a safety model to define a background in which the system 100 can perform autonomous control of the vehicle 200 in a safety mode. In order to develop a safety model, in some embodiments, the processing unit 110 may consider the positioning and movement of other vehicles, detected road edges and obstacles, and/or extracted from map data (such as data from the map database 160) A description of the rough road shape. By considering additional information sources, the processing unit 110 can provide redundancy for detecting road markings and lane geometry and increase the reliability of the system 100.

[0121]

Figure 5D is a flowchart showing an exemplary process 500D for detecting traffic lights in a set of images consistent with the disclosed embodiments.

The processing unit 110 may cause the monocular image analysis module 402 to execute the process 500D. In step 560, the processing unit 110 may scan the set of images and identify objects appearing at locations in the images that may include traffic lights. For example, the processing unit 110 may filter the recognized objects to construct a set of candidate objects that exclude those objects that are unlikely to correspond to traffic lights. Filtering can be done based on various attributes associated with the traffic light, such as shape, size, texture, positioning (relative to the vehicle 200, etc.). Such attributes can be based on multiple examples of traffic lights and traffic control signals, and can be stored in a database. In some embodiments, the processing unit 110 may perform multi-frame analysis on the set of candidate objects reflecting possible traffic lights. For example, the processing unit 110 may track candidate objects in consecutive image frames, estimate the true location of the candidate objects, and filter out those moving objects (which cannot be traffic lights). In some embodiments, the processing unit 110 may perform color analysis on the candidate object, and recognize the relative location of the detected color that appears in a possible traffic light.

[0122]

In step 562, the processing unit 110 may analyze the geometry of the intersection.

The analysis can be based on any combination of: (i) the number of lanes detected on any side of the vehicle 200, (ii) the marks (such as arrow marks) detected on the road, and (iii) from the map data The description of the extracted intersection (such as data from the map database 160). The processing unit 110 may use the information obtained by executing the monocular analysis module 402 to perform the analysis. In addition, the processing unit 110 may determine the correspondence between the traffic light detected in step 560 and the lane appearing near the vehicle 200.

[0123]

In step 564, when the vehicle 200 approaches the intersection, the processing unit 110 may update the confidence level associated with the analyzed intersection geometry and the detected traffic light.

For example, compared to the number of traffic lights that actually appear at an intersection, it is estimated that the number of traffic lights that appear at an intersection may affect the confidence level. Therefore, based on the confidence level, the processing unit 110 may delegate control to the driver of the vehicle 200 in order to improve safety conditions. By performing

IPR2025-00943
Tesla EX1002 Page 1060

steps 560-564, the processing unit 110 can recognize the traffic lights appearing in the group of captured images, and analyze the information of the geometric shape of the intersection. As described above in conjunction with FIG. 5A, based on recognition and analysis, the processing unit 110 may cause one or more navigation responses in the vehicle 200.

[0124]

5E is a flowchart illustrating an exemplary process 500E for eliciting one or more navigation responses in the vehicle 200 based on the vehicle path consistent with the disclosed embodiments.

In step 570, the processing unit 110 may construct an initial vehicle path associated with the vehicle 200. The vehicle path may be represented by a set of points expressed in coordinates (x, z), and in some embodiments, the distance di between two points in the set of points may fall within the range of 1 meter to 5 meters, Or in other embodiments, timing measurements may be used to provide the set of points, where the points in the set of points are separated from each other by 0.5 seconds to 5 seconds (e.g., 1.5 seconds apart). In one embodiment, the processing unit 110 may use two polynomials (such as polynomials for the left side road and the right side road) to construct the initial vehicle path. If any (zero offset may correspond to the travel in the middle of the lane), the processing unit 110 may calculate the geometric midpoint between the two polynomials, and offset each point included in the resulting vehicle path by a predetermined Offset (for example, smart lane offset). The offset can be in a direction perpendicular to the segment between any two points in the vehicle path. In another embodiment, the processing unit 110 may use a polynomial and the estimated lane width to offset each point of the vehicle path by half the estimated lane width plus a predetermined offset (for example, smart lane offset).).

[0125]

In step 572, the processing unit 110 may update the vehicle path constructed in step 570.

The processing unit 110 may reconstruct the vehicle path constructed in step 570 with a higher resolution, so that the distance dk between two points in the group of points representing the vehicle path is smaller than the distance di described above. For example, the distance dk may fall within a range of 0.1 meters to 0.3 meters. The processing unit 110 may use a parabolic spline algorithm to reconstruct the vehicle path, which may derive a cumulative distance vector S corresponding to the total length of the vehicle path (ie, based on the set of points representing the vehicle path).

[0126]

In step 574, the processing unit 110 may determine a look-ahead point (expressed as (x1, z1) in coordinates) based on the updated vehicle path constructed in step 572.

The processing unit 110 may extract an advance point from the accumulated distance vector S, and the advance point may be associated with an advance distance and an advance time. The advance distance, which may have a lower limit in the range of 10 meters to 20 meters, may be calculated as the product of the speed of the vehicle 200 and the advance time. For example, as the speed of the vehicle 200 decreases, the leading distance may also decrease (for example, until the leading distance reaches a lower limit). The lead time, which can range from 0.5 seconds to 1.5 seconds, can be inversely proportional to the gain of one or more control loops (such as heading error tracking control loops) that are related to causing a navigation response in the vehicle 200 United. For example, the gain of the heading error tracking control loop may depend on the vehicle's lateral power, the steering actuator loop, the bandwidth of the yaw rate loop, and so on. Therefore, the higher the gain of the heading error tracking control loop, the lower the lead time.

[0127]

In step 576, the processing unit 110 may determine a heading error and a yaw rate command based on the advance point determined in step 574.

The processing unit 110 may determine the heading error by calculating the arctangent of the leading point (for example, arctan (x1/z1)). The processing unit 110 may determine that the yaw rate command is the product of the heading error and the high-level control gain. The high level control gain can be equal to: (2/advance time), if the advancing distance is not above the lower limit. Otherwise, the high-level control gain can be equal to: (2*vehicle 200 speed/leading distance).

[0128]

IPR2025-00943
Tesla EX1002 Page 1061

FIG. 5F is a flowchart showing an exemplary process 500F for determining whether a leading vehicle is changing lanes, consistent with the disclosed embodiment.

In step 580, the processing unit 110 may determine navigation information associated with a leading vehicle (e.g., a vehicle traveling in front of the vehicle 200). For example, the processing unit 110 may use the techniques described above in conjunction with FIGS. 5A and 5B to determine the position, speed (for example, direction and velocity) and/or acceleration of the leading vehicle. The processing unit 110 may also use the technique described above in conjunction with FIG. 5E to determine one or more road polynomials, advance points (associated with the vehicle 200) and/or snail trails (for example, describing the leading vehicle A set of points of the path taken).

[0129]

In step 582, the processing unit 110 may analyze the navigation information determined in step 580.

In one embodiment, the processing unit 110 may calculate the distance between the snail trail and the path polynomial (for example, along the trail). If the variance of the distance along the trail exceeds a predetermined threshold (for example, 0.1 m to 0.2 m on a straight road, 0.3 m to 0.4 m on a moderately curved road, and 0.5 m to 0.6 m on a road with sharp curves), the processing unit 110 may determine that the leading vehicle may be changing lanes. In the event that multiple vehicles are detected to be traveling in front of the vehicle 200, the processing unit 110 may compare the worming trails associated with each vehicle. Based on the comparison, the processing unit 110 may determine that a vehicle whose snail trail does not match the snail trail of other vehicles may be changing lanes. In addition, the processing unit 110 may compare the curvature of the snail track (associated with the leading vehicle) with the expected curvature of the road segment traveled by the leading vehicle. The expected curvature can be extracted from map data (for example, data from the map database 160), road polynomials, snail track of other vehicles, prior knowledge about the road, and so on. If the difference between the snail track and the curvature and the expected curvature of the road segment exceeds a predetermined threshold, the processing unit 110 may determine that the leading vehicle may be changing lanes.

[0130]

In another embodiment, the processing unit 110 may compare the instantaneous positioning of the leading vehicle with the leading point (associated with the vehicle 200) within a designated time period (for example, 0.5 seconds to 1.5 seconds).

If during the specified time period, the distance between the instantaneous positioning of the leading vehicle and the leading point changes, and the cumulative change exceeds a predetermined threshold (for example, 0.3 meters to 0.4 meters on a straight road, 0.7 meters on a moderately curved road To 0.8 meters, and 1.3 meters to 1.7 meters on a road with a sharp curve), the processing unit 110 may determine that the leading vehicle may be changing lanes. In another embodiment, the processing unit 110 may analyze the geometry of the snail trail by comparing the lateral distance traveled along the trail with the expected curvature of the snail trail. The desired radius of curvature can be determined according to the following calculation: $(\delta z^2+\delta x^2)/2/(\delta x)$, where δx represents the lateral distance traveled, and δz represents the longitudinal distance traveled. If the difference between the lateral distance traveled and the desired curvature exceeds a predetermined threshold (for example, 500 meters to 700 meters), the processing unit 110 may determine that the leading vehicle may be changing lanes. In another embodiment, the processing unit 110 may analyze the location of the leading vehicle. If the positioning of the leading vehicle obscures the road polynomial (eg, the leading vehicle is overlapped on top of the road polynomial), the processing unit 110 may determine that the leading vehicle may be changing lanes. In a situation where the positioning of the leading vehicle is detecting that another vehicle is in front of the leading vehicle and the snail trails of the two vehicles are not parallel, the processing unit 110 may determine that the (closer) leading vehicle may be on track. Change lanes.

[0131]

In step 584, the processing unit 110 may determine whether the leading vehicle 200 is changing lanes based on the analysis performed in step 582.

For example, the processing unit 110 may make the determination based on the weighted average of the various analyses performed in step 582. Under such a scheme, for example, the processing unit 110's determination that the leading vehicle may be changing lanes based on a specific type of analysis may be assigned the value "1" (and "0" indicates that the leading vehicle is not like the determination that the leading vehicle is changing lanes) . The different analyses performed in step 582 may be assigned different weights, and the disclosed embodiments are not limited to any specific combination of analysis and weights.

[0132]

FIG. 6 is a flowchart showing an exemplary process 600 for eliciting one or more navigation responses based on stereoscopic image analysis consistent with the disclosed embodiments.

In step 610, the processing unit 110 may receive the first plurality of images and the second plurality of images via the data interface 128. For example, a camera included in the image capturing unit 120 (such as image capturing devices 122 and 124 having fields of view 202 and 204) can capture the first and second plurality of images of the forward area of the vehicle 200, and They are transmitted to the processing unit 110 through a digital connection (for example, USB, wireless, Bluetooth, etc.). In some embodiments, the processing unit 110 may receive the first plurality of images and the second plurality of images via two or more data interfaces. The disclosed embodiments are not limited to any specific data interface configuration or protocol.

[0133]

In step 620, the processing unit 110 may cause the stereo image analysis module 404 to perform stereo image analysis on the first plurality of images and the second plurality of images to create a 3D map of the road in front of the vehicle and detect features in the images, such as Lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, road dangerous objects, etc.

The stereoscopic image analysis can be performed in a manner similar to the steps described above in conjunction with FIGS. 5A-5D. For example, the processing unit 110 may enable the stereoscopic image analysis module 404 to perform detection of candidate objects (for example, vehicles, pedestrians, road markings, traffic lights, road hazards, etc.) in the first plurality of images and the second plurality of images. Filter out a subset of candidate objects based on various criteria, perform multi-frame analysis, construct measurements, and determine confidence levels about the remaining candidates. When performing the above steps, the processing unit 110 may consider information from both the first plurality of images and the second plurality of images, instead of information from a group of images alone. For example, the processing unit 110 may analyze differences in pixel-level data (or other subsets of data from the two captured image streams) for candidate objects appearing in both the first plurality of images and the second plurality of images. Another example is that the processing unit 110 may observe that the object appears in one of the multiple images but does not appear in the other images, or the object may be different from the objects appearing in the two image streams. To estimate the location and/or speed of the candidate object (eg, relative to the vehicle 200). For example, the location, speed, and/or acceleration relative to the vehicle 200 may be determined based on the trajectory, location, movement characteristics, etc. of features associated with objects appearing in one or two image streams.

[0134]

In step 630, the processing unit 110 may cause the navigation response module 408 to perform one or more navigation responses in the vehicle 200 based on the analysis performed in step 620 and the technique described above in conjunction with FIG. 4.

For example, the navigation response may include turns, lane shifts, changes in acceleration, changes in speed, braking, etc. In some embodiments, the processing unit 110 may use data derived from the execution of the speed and acceleration module 406 to cause one or more navigation responses. In addition, multiple navigation responses can occur simultaneously, sequentially, or any combination of them.

[0135]

FIG. 7 is a flowchart showing an exemplary process 700 for eliciting one or more navigation responses based on analysis of three sets of images consistent with the disclosed embodiments.

In step 710, the processing unit 110 may receive the first plurality of images, the second plurality of images, and the third plurality of images via the data interface 128. For example, a camera included in the image capture unit 120 (such as image capture devices 122, 124, and 126 having fields of view 202, 204, and 206) may capture the first plurality of images of the forward area and/or sides of the vehicle 200 , The second plurality of images and the third plurality of images, and transmit them to the processing unit 110 through a digital connection (for example, USB, wireless, Bluetooth, etc.). In some embodiments, the processing unit 110 may receive the first plurality of images, the second plurality of images, and the third plurality of images via three or more data interfaces. For example, each of the image capturing devices 122, 124, and 126 may have an associated

data interface for transferring data to the processing unit 110. The disclosed embodiments are not limited to any specific data interface configuration or protocol.

[0136]

In step 720, the processing unit 110 may analyze the first plurality of images, the second plurality of images, and the third plurality of images to detect features in the images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, Traffic lights, road dangerous objects, etc.

The analysis can be performed in a manner similar to the steps described above in conjunction with FIGS. 5A-5D and FIG. 6. For example, the processing unit 110 may perform processing on each of the first plurality of images, the second plurality of images, and the third plurality of images (for example, by executing the monocular image analysis module 402 and based on the above described in conjunction with FIGS. 5A-5D). Step) Perform monocular image analysis. Alternatively, the processing unit 110 may compare the first plurality of images and the second plurality of images, the second plurality of images and the third plurality of images, and/or the first plurality of images and the third plurality of images (for example, The stereoscopic image analysis is performed by executing the monocular image analysis module 404 and based on the steps described above in conjunction with FIG. 6). The information processed by the analysis corresponding to the first plurality of images, the second plurality of images, and/or the third plurality of images may be combined. In some embodiments, the processing unit 110 may perform a combination of monocular image analysis and stereo image analysis. For example, the processing unit 110 may perform monocular image analysis on the first plurality of images (for example, by executing the monocular image analysis module 402), and perform monocular image analysis on the second pluraiity of images and the third piurality of images (for example, by executing the stereoscopic image The analysis module 404) performs stereo image analysis. The configuration of the image capture devices 122, 124, and 126----including their respective positions and fields of view 202, 204, and 206----can affect the analysis of the first, second, and third images. type. The disclosed embodiments are not limited to the specific configuration of the image capture devices 122, 124, and 126 or the type of analysis performed on the first plurality of images, the second plurality of images, and the third plurality of images.

[0137]

In some embodiments, the processing unit 110 may perform a test on the system 100 based on the images collected and analyzed in steps 710 and 720.

Such testing may provide an indicator of the overall performance of the system 100 for certain configurations of image capture devices 122, 124, and 126. For example, the processing unit 110 may determine the ratio of "false hits" (for example, a situation where the system 100 does not correctly determine the presence of vehicles or pedestrians) and "missings".

[0138]

In step 730, the processing unit 110 may cause one or more navigation responses in the vehicle 200 based on information derived from two of the first plurality of images, the second plurality of images, and the third plurality of images.

The selection of two of the first plurality of images, the second plurality of images, and the third plurality of images may depend on various factors, such as, for example, the number and type of objects detected in each of the plurality of images And size. The processing unit 110 may also be based on the quality and resolution of the image, the effective field of view reflected in the image, the number of captured frames, and the extent to which one or more objects of interest actually appear in the frame (for example, the object in which the object appears) The percentage of frames, the proportion of objects appearing in each such frame, etc.) etc. to make choices.

[0139]

In some embodiments, the processing unit 110 may select from the first plurality of images, the second plurality of images, and the third plurality of images by determining the degree to which the information derived from one image source is consistent with the information derived from other image sources. Information derived from two of the images.

For example, the processing unit 110 may combine (through monocular analysis, stereo analysis, or any combination of the two) the processed information derived from each of the image capture devices 122, 124, and 126, and determine whether the information is received from the image capture device Consistent visual indicators (eg, lane markings, detected vehicles and their locations and/or paths, detected traffic lights, etc.) in each of the captured images of 122, 124, and 126. The processing

unit 110 may also exclude information that is inconsistent in the captured image (for example, a vehicle that changes lanes, a lane model indicating that the vehicle is too close to the vehicle 200, etc.). Therefore, the processing unit 110 may select information derived from two of the first plurality of images, the second plurality of images, and the third plurality of images based on the consistent determination and the inconsistent information.

[0140]

For example, the navigation response may include turns, lane shifts, changes in acceleration, and so on.

The processing unit 110 may cause one or more navigation responses based on the analysis performed in step 720 and the technique described above in conjunction with FIG. 4. The processing unit 110 may also use the data obtained by executing the speed and acceleration module 406 to cause one or more navigation responses. In some embodiments, the processing unit 110 may be based on the relative positioning, relative velocity and/ Or relative acceleration to cause one or more navigation responses. Multiple navigation responses can occur simultaneously, sequentially, or any combination of them.

[0141]

Lane restriction prediction

[0142]

The system 100 can provide the vehicle 200 with lane restriction prediction technology, so that the system predicts the location of one or more lane restrictions.

[0143]

In some embodiments, lane restriction prediction may be understood as one or more techniques for predicting the location of lane restrictions.

In some embodiments, a lane restriction may be defined as a boundary line at the edge of a lane, road, parking space, or other path or space that enables vehicles to travel.

In some cases, lane restrictions may be marked on the road, such as through drawn lines, physical obstacles, bots points, or one or more additional indicators.

In some embodiments, lane restrictions may be defined by the edges of roads, curbs, guardrails, and the like. In some embodiments, the lane restriction may be defined by one or more other vehicles (such as vehicles parked in parallel rows in adjacent lanes, or vehicles parked in rows in a parking lot). In some embodiments, lane restrictions may not be marked, such as on unmarked roads or in parking lots without lane lines.

[0144]

In some embodiments, the self-driving system or driver assistance system may use image capture devices and image analysis to determine one or more objects based on detected objects in the image (such as lane markings, road edges, and/or other vehicles). The location of the lane restriction.

In most driving situations, the system may be able to detect at least two lane restrictions based on the visible characteristics of the captured image (for example, lane lines, etc.)-one on each side of the lane the vehicle travels.

[0145]

However, in some cases, only one lane restriction included on one side of the lane or on the side of the vehicle 200 may be sufficiently detected based on the visible characteristics (for example, lane lines, etc.) in the captured image. In those cases, the system 100 can predict the location of the lane restriction on the other side of the lane or on the other side of the vehicle, enabling the vehicle to continue in the lane defined by the detected restriction and the predicted restriction against it March.

[0146]

IPR2025-00943
Tesla EX1002 Page 1065

In some embodiments, the lane restriction prediction may be performed by the system 100 premised on satisfying one or more conditions.

For example, in some embodiments, the system 100 predicts the location of the lane restriction when the previously detected lane restriction is stopped as the vehicle 200 advances along the road.

Therefore, in some embodiments, if lane markings or other detected lane restrictions suddenly end or become undetectable, lane restriction prediction may be performed. In other examples, when the system determines that the guardrail shadow condition exists, the system 100 predicts the location of the lane restriction so that the guardrail exists near the road or the guardrail casts a shadow on the road.

[0147]

In some embodiments, when the system 100 has access to valid and sufficient historical information about the lane width traveled by the vehicle 200, the system 100 may only perform lane restriction prediction; the lane restriction on one side of the vehicle 200 follows the vehicle 200 In the case where 200 is newly detected along the road, in some embodiments, the system 100 may not perform lane restriction prediction.

[0148]

In some embodiments, the lane restriction prediction may be performed within the maximum travel distance, so that if the vehicle 200 travels a distance greater than the predetermined maximum travel distance while the system 100 continuously performs the lane restriction prediction, the lane restriction prediction may subsequently be stop.

In some embodiments, the predicted maximum travel distance may be set or modified according to one or more various factors detected by the system 100.

In some embodiments, the predicted maximum travel distance may be turned according to one or more various factors detected by the system 100 (for example, it may be functionally set to "infinite").

[0149]

In some embodiments, the lane restriction prediction may be stopped according to various cancellation conditions detected by the system 100.

For example, in some embodiments, if the driver of the vehicle 200 changes lanes, the lane restriction prediction may stop. In some embodiments, if the vehicle 200 traverses a curve in the road that exceeds a predetermined curve sharpness threshold, the lane restriction prediction may stop.

[0150]

8A-8C illustrate an embodiment in which the vehicle 200 travels on the road 800. As shown in FIG.

In some embodiments, the vehicle 200 may have one or more fields of view provided by one or more image capture devices or cameras, such as image capture devices 122, 124, and/or 126. In the depicted example, as depicted by an angled line forming an upwardly opening triangle in front of the vehicle 200, the vehicle 200 may have a field of view from the vehicle toward the front; in some embodiments, the vehicle 200 may have The vehicle 200 faces one or more fields of view in a plurality of different directions. In the depicted example, the disclosed systems and systems for lane restriction prediction (for example, predicting the location of the lane restriction when no lane restriction is detected) and operating the vehicle 200 in the lane defined by the predicted lane restriction may be used. method. As shown, the road 800 may have a lane 802 in which the vehicle 200 travels. The lane 802 may be bounded by a first lane restriction 804 on a first side, and a lane restriction 806 on a second side opposite the first side. As indicated by the bold lines showing lane restrictions 804 and 806, these lane restrictions 804 and 806 can be detected by the system 100. The vehicle 200 may have a first vehicle side 808 on the same side of the vehicle 200 as the first lane restriction 804 and a second vehicle side 810 on the same side of the vehicle 200 as the second lane restriction 806.

[0151]

IPR2025-00943
Tesla EX1002 Page 1066

The processing unit 110 may be configured to determine the first lane restriction 804 and the second lane restriction 806 based on a plurality of images captured by the image capturing devices 122-126, which may be received by the processing unit 110 via the data interface 128.

According to some embodiments, the first lane restriction 804 and/or the second lane restriction 806 may be visible lane demarcation lines (such as dashed or solid lines marked on the road surface like lane lines 803a (solid lines) and 803b (dashed lines)) identify. Additionally or alternatively, the first lane restriction 804 and/or the second lane restriction 806 may be identifiable by obstacles or edges of the road surface. Additionally or alternatively, the first lane restriction 804 and/or the second lane restriction 806 may be identified by markings (e.g., Botts point). According to some embodiments, the processing unit 110 may determine the first lane restriction 804 and/or the second lane restriction 806 by identifying the midpoint of the road surface width. According to some embodiments, if the processing unit 110 recognizes only the first lane restriction 804, the processing unit 110 may estimate the second lane restriction 806, such as based on the estimated lane width or road width. For example, when the line of the designated road lane is not drawn or marked by other means, the processing unit 110 can predict the lane restriction in this way.

[0152]

In FIGS. 8A-8C, lane restrictions 804 and 806 are marked on the road 800 by solid lines and dashed lines 803a and 803b, respectively.

In the portion of the road 800 toward the center of FIGS. 8A to 8C, the broken line 803 b corresponding to the lane restriction 806 is not drawn on the road 800. In FIGS. 8A to 8C, the detection of lane restrictions 804 and 806 by the system 100 is depicted by the bold lines covering the corresponding lane lines in front of the vehicle 200. As depicted by the bold line in FIG. 8A, the system 100 detects a lane restriction 804 as a solid lane restriction on the left side of the vehicle 200 within a certain visible distance in front of the vehicle 200. Similarly, the system 100 detects a lane restriction 806 as a dashed lane restriction on the right side of the vehicle 200 within a certain visible distance in front of the vehicle 200. Based on the detection of the lane restrictions 804 and 806, the system 100 causes the vehicle 200 to travel in the lane 802 defined by the lane restrictions 804 and 806.

[0153]

FIG. 8B depicts the vehicle 200 continuing along the road 800 in the lane 802 after the point in time depicted in FIG. 8A.

When the vehicle 200 reaches a portion of the road 800 where the dashed line 803b is not drawn on the road 800, the lane restriction prediction may be performed by the system 100. As shown, the system 100 predicts that the location of the lane restriction 806 is at a distance 812 from the opposite lane restriction 804. In some embodiments, the distance 812 may be determined based on the detected historical information and the previous width of the lane 802 stored by the system 100 and/or the approximate average historical width of the lane. As shown in the figure, the system 100 predicts that the lane restriction 806 will continue to be a virtual lane restriction based on the fact that the lane restriction 806 was detected by the system 100 for the last time the lane restriction 806 was marked by the virtual lane line 803b before the lane restriction prediction was activated. Based on the continued detection of the lane restriction 804 and the prediction of the lane restriction 806, the system 100 causes the vehicle 200 to continue traveling in the lane 802 defined by the lane restrictions 804 and 806. In some embodiments, when the lane restriction prediction is activated, and possibly under certain circumstances (for example, there are other vehicles nearby), a certain default type of lane restriction is used to control the movement of the vehicle, and the default type of lane restriction is not Must be the same as the last lane restriction detected before the lane restriction prediction was activated. For example, in some embodiments, the system may default to predicting solid lane restrictions instead of dashed lane restrictions.

[0154]

FIG. 8C depicts the vehicle 200 continuing along the road 800 in the lane 802 after the point in time depicted in FIG. 8B.

When the vehicle 200 reaches the part of the road 800 drawn on the road 800 again by the dashed line 803b, the lane restriction 806 is again detected by the system 100 as a virtual lane restriction (rather than being predicted), so the lane restriction prediction can be made by the system 100 stop execution. As depicted in FIG. 8C, the vehicle 200 traveled a distance 814 when performing the lane restriction prediction. In some embodiments, lane restriction prediction may only be performed continuously or substantially continuously within a predetermined maximum distance (and/or predetermined maximum time). In the example depicted in FIG. 8C, the distance 814 is less than the predetermined maximum distance used to perform the lane restriction prediction, so that the lane restriction prediction is different from its location in FIG. 8B as the vehicle 200 traverses its

IPR2025-00943
Tesla EX1002 Page 1067

location in FIG. 8B and its location in FIG. 8C. The distance between can be executed continuously. In some embodiments, the predetermined maximum distance for performing lane restriction prediction may be set to a default value of 18 meters. In some embodiments, the predetermined maximum distance may be changed by the vehicle manufacturer, driver, and/or any other personnel.

[0155]

9A-9C illustrate an embodiment in which the vehicle 200 is traveling on a road 900. As shown in FIG.

In the depicted example, the disclosed systems and systems for lane restriction prediction (for example, predicting the location of the lane restriction when no lane restriction is detected) and operating the vehicle 200 in the lane defined by the predicted lane restriction may be used. method. The example depicted in FIGS. 9A-9C is different from the example depicted in FIGS. 8A-8C at least in that two lanes on the road 900 are merged into one lane. As explained in more detail below, in some embodiments, when two lanes are merged into one lane and the lane line is not detected on the side of the lane that merges the second lane, lane restriction prediction may be performed.

[0156]

As shown, the road 900 may have a lane 902 in which the vehicle 200 travels.

The lane 902 may be bounded by a first lane restriction 904 on a first side, and by a lane restriction 906 on a second side opposite the first side. As indicated by the bold lines showing lane restrictions 904 and 906, these lane restrictions 904 and 906 can be detected by the system 100. The vehicle 200 may have a first vehicle side 908 on the same side of the vehicle 200 as the first lane restriction 904 and a second vehicle side 910 on the same side of the vehicle 200 as the second lane restriction 906.

[0157]

The processing unit 110 may be configured to determine the first lane restriction 904 and the second lane restriction 906 based on a plurality of images captured by the image capturing devices 122-126, which may be received by the processing unit 110 via the data interface 128.

According to some embodiments, the first lane restriction 904 and/or the second lane restriction 906 can be identified, detected, estimated, or predicted in any or all of the ways discussed above with reference to FIGS. 8A-8C.

[0158]

As shown, the road 900 may have a lane 916.

The lane 916 may be restricted by the lane. In the example shown, the lane 916 is marked by a solid line 918 on its left side, and a solid line 120 on its right side. In the illustrated example, the lane 916 is merged into the lane 902 as a road 900 that continues in the direction of travel of the vehicle 200.

[0159]

In FIGS. 9A-9C, the lane restriction 904 is continuously marked by a solid line 903 on the road 900.

In addition, the lane restriction 906 is marked inconsistently; in the bottom part of the figure, the lane restriction 906 is marked by a solid line 905a on the road 900; in the part above the bottom part, as the lane 916 starts to merge into the lane 902, The lane restriction 906 is marked by the dashed line 905b on the road 900; in the part above this part, as the lane 916 is completed and merged into the lane 902, the lane restriction 906 is not marked; finally, in the topmost part of the figure, the lane restriction 906 It is again marked by a solid line, and this time on the road 900 by a solid line 905c.

[0160]

In FIGS. 9A-9C, the detection of lane restrictions 904, 906, and 922 by the system 100 is depicted by bold lines covering the corresponding lane lines in front of the vehicle 200.

IPR2025-00943
Tesla EX1002 Page 1068

As depicted by the bold line in FIG. 9A, the system 100 detects a lane restriction 904 as a solid lane restriction on the left side of the vehicle 200 within a certain visible distance in front of the vehicle 200. Similarly, the system 100 detects the lane restriction 906 as the actual lane restriction on the right side of the vehicle 200 within a certain visible distance in front of the vehicle 200. Based on the detection of the lane restrictions 904 and 906, the system 100 causes the vehicle 200 to travel in the lane 902 defined by the lane restrictions 904 and 906.

[0161]

FIG. 9B depicts the vehicle 200 continuing along the road 900 in the lane 902 after the point in time depicted in FIG. 9A.

As the vehicle 200 reaches the portion of the road 900 where the lane 916 merges from the right into the lane 902, where the lane restriction 906 is briefly marked by the dashed line 905b and then not marked by any line, the lane restriction prediction may be performed by the system 100. As shown, the system 100 predicts that the location of the lane restriction 906 is at a distance 912 from the opposite lane restriction 904. In some embodiments, the distance 912 may be determined based on the detected historical information and the previous width of the lane 902 stored by the system 100 and/or the approximate average historical width of the lane. As shown in the figure, the system 100 predicts that the lane restriction 906 continues to be a virtual lane restriction based on the fact that the lane restriction 906 was detected by the system 100 the last time the lane restriction 906 was marked by the dashed line 905b before the lane restriction prediction was activated. Based on the continued detection of the lane restriction 904 and the prediction of the lane restriction 906, the system 100 causes the vehicle 200 to continue traveling in the lane 902 defined by the lane restrictions 904 and 906.

[0162]

In addition, shown in FIG. 9B is a system 100 that detects a lane restriction 922 that merges with the lane 916 to the right side of the lane 902 that bounds the lane 916.

The system 100 detects the lane restriction 922 as the actual lane restriction based on the corresponding solid line 920 on the road 900. In some embodiments, the system 100 may determine that the detected lane restriction 922 is not a restriction of the lane 902 in which the vehicle 200 travels, but a lane restriction of a lane incorporated into the lane 902. In some embodiments, this determination may be made based on a determination that the lane restriction 922 is detected to approach the lane 902 above the threshold rate of approach. For example, if the detected lane restriction outside the restriction of the lane in which the vehicle is traveling makes an angle greater than a predetermined threshold angle or more parallel to the detected or previously detected lane restriction of the lane in which the vehicle is traveling, the outer lane can be determined The restriction corresponds to the approaching merged lane.

[0163]

In some embodiments, the determination of lane merging into another lane may be made based on detection and analysis of other aspects, characteristics, or objects in the vehicle or the road around it.

For example, a second vehicle (e.g., different from the vehicle providing the system) may be detected and tracked, and the system may determine that another vehicle is traveling at an angle to the lane on the road, that is, the lane instructing the other vehicle to merge into In another lane. In some embodiments, the system may detect that another vehicle approaches the lane from an area outside the lane above the threshold rate of approach. In some embodiments, the system may make such a determination regarding one or more other vehicles, and may accordingly determine lane merging around the one or more other vehicles. In some embodiments, information about one or more other vehicles may be received through multiple images captured by a camera associated with the system; in some embodiments, alternatively or additionally, information about one or more The information of the other vehicles may be received through any suitable form of electronic communication, such as communication transmitted from one or more of the one or more other vehicles.

[0164]

In some other embodiments, the determination of lane merging may be made based on historical information about lane configuration and/or layout (such as historical information of lanes at or around the current position of the vehicle).

[0165]

IPR2025-00943
Tesla EX1002 Page 1069

Figure 9C depicts the vehicle 200 continuing along the road 900 in the lane 902 after Figure 9B.

When the vehicle 200 reaches the portion of the road 900 that is merged into the lane 902 in the lane 916, the lane restriction 906 is again detected by the system 100 as a virtual lane restriction (rather than being predicted), so the lane restriction 906 is again The road 900 is marked by a solid line, this time by a solid line 905c, and the lane restriction prediction may stop being executed by the system 100.

In some embodiments, the system 100 may determine the lane limit of the lane in which the vehicle 200 is detected by determining that the lane limit is no longer close to a vehicle (or the lane in which the vehicle travels) above a predetermined approaching threshold rate (rather than a combined Lane restrictions). In some embodiments, the system 100 may determine the lane limit of the lane in which the vehicle is detected (instead of the combined lane) by determining that the lane limit is within a predetermined threshold distance of the predicted position of the lane limit as predicted by the lane limit prediction. Lane restrictions).

[0166]

As depicted in FIG. 9C, the vehicle 200 traveled a distance 914 when performing the lane restriction prediction.

In some embodiments, lane restriction prediction may only be performed continuously or substantially continuously within a predetermined maximum distance (and/or predetermined maximum time). According to some embodiments, the maximum distance/time allocated for lane restriction prediction may be dynamically set. For example, the maximum distance/time allocated for lane restriction prediction can be based on the surrounding environment of the vehicle, the speed of the vehicle, the topology of the road, the brightness or environment, the weather, the density and/or proximity of other vehicles on the road, and/ Or any other road conditions or characteristics to adjust. In some embodiments, when it is determined that a close merged lane exists in the vicinity of the vehicle 200, the system 100 may permit the lane restriction prediction to be performed within a larger total distance than the lane restriction prediction should otherwise be permitted to be performed. For example, although the default maximum distance may be 18 meters, in the case of detecting an ongoing lane merging, the maximum distance may be 1000 meters. This difference may be due to consideration of the fact that lane merging may require several hundred meters to complete. In the example depicted in FIG. 8C, the distance 914 is less than the predetermined maximum distance used to perform lane restriction prediction in the case of lane merging, so that the lane restriction prediction can be as the vehicle 200 traverses its positioning in FIG. 9B and The distance between its positioning in FIG. 9C is continuously performed.

[0167]

10A and 10B illustrate an embodiment in which the vehicle 200 travels on the road 1000.

FIG. 10A shows a three-dimensional view as seen from the perspective of the vehicle 200, such as one of the one or more images received by the system 100 when the vehicle 200 is traveling on the road 1000. (FIG. 10A depicts the leading vehicle driving on the road 1000 in front of the vehicle 200, not to be confused with the vehicle 200 itself.) FIG. 10B shows an overhead view of the vehicle 200 on the same part of the road 1000.

[0168]

In the depicted example, the disclosed systems and methods for lane restriction prediction (for example, predicting the location of the lane restriction when a guardrail shadow condition exists) and operating the vehicle 200 in the lane defined by the predicted lane restriction may be used .

As shown in the figure, the road 1000 may have a lane 1002 in which the vehicle 200 travels. The lane 1002 may be bounded by a first lane restriction 1004 on a first side, and by a lane restriction 1006 on a second side opposite to the first side. In some embodiments, as indicated by the bold lines showing lane restrictions 1004 and 1006, these lane restrictions 1004 and 1006 may be detected by the system 100. The vehicle 200 may have a first vehicle side 1008 on the same side of the vehicle 200 as the first lane restriction 1004, and a second vehicle side 1010 of the vehicle 200 on the same side as the second lane restriction 1006.

[0169]

The road 1000 may include lane lines 1003a and 1003b, which physically mark the left and right sides of the lane 1002 on the road 1000, respectively.

IPR2025-00943
Tesla EX1002 Page 1070

[0170]

The road 1000 may include a guardrail 1012 supported by a pillar 1014.

Both the guardrail 1012 and its pillar 1014 cast a guardrail shadow 1016, which in the depicted example falls on the lane line 1003a and into the lane 1002.

[0171]

As depicted by the bold lines in FIGS. 10A and 10B, the system 100 detects a lane restriction 1006 as a solid line restriction on the right side of the vehicle 200 within a certain visible distance in front of the vehicle 200.

On the left side of the road 1000, the bold line shows how the system 1000 detects some edges, one of which is attributable to lane line 1003a, and two of them (shaded edges 1005a and 1005b) are attributable to guardrail shadow 1016 .

In some embodiments, the detection of multiple edges on one side of a vehicle or lane may make it more difficult to recognize lane restrictions from images available to the system 100, so that lane restrictions cannot be recognized with a high level of certainty. Therefore, in some embodiments, when the system 100 determines that the guardrail 1012 or the guardrail post 1014 exists, that is, the guardrail shadow 1016 is projected on the road 1000, or multiple edges are detected on the same side of the lane 1002, the system 100 It can be determined that the guardrail shadow condition exists, and the guardrail shadow mode can be entered accordingly. In some embodiments, the guardrail shadow mode enables lane restriction predictions to be performed on the side of the lane approaching the guardrail, shadow, or edge being detected. By using lane restriction prediction, the system 100 can predict the location of the lane restriction 1004 according to any of the techniques discussed above with reference to FIGS. 8A-8C and/or FIGS. 9A-9C. In the depicted example, the system 100 can correctly predict the location of the lane restriction 1002 that coincides with the edge defined by the drawn lane line 1003a instead of the edge defined by the guardrail shadow 1016. Based on the detection of the lane restrictions 1004 and 1006, the system 100 causes the vehicle 200 to travel in the lane 1002 defined by the lane restrictions 1004 and 1006.

[0172]

In some embodiments, guardrail shadow prediction may allow lane restriction prediction to be performed without restrictions on how long the lane restriction prediction may be performed.

Therefore, for example, the maximum distance that the system 100 can bend to instead of the maximum distance for lane restriction prediction is set to 18 meters (in the example of FIGS. 8A-8C) or 1000 meters (in the example of FIGS. 9A-9C) , Functionally configure the system so that as long as the guardrail shadow condition continues to be detected and as long as the lane restriction prediction cancellation condition is not detected, the lane restriction prediction can continue within an infinite distance (and/or infinite time).

[0173]

FIG. 11 is an exemplary functional block diagram of the memory 140 or 150, which can store instructions for performing one or more operations consistent with the disclosed embodiments.

As illustrated in FIG. 11, the memory 140 or 150 may store one or more modules for performing the techniques described herein. For example, the memory 140 or 150 may store an image acquisition module 1100, a road lane restriction recognition module 1102, a determination module 1104, a road lane restriction prediction module 1106, a vehicle start module 1108, and a lane restriction prediction stop module 1110. As those of ordinary skill in the art will recognize, any one of the modules listed above may be combined with each other module or other modules, may be combined with another module or other modules, and/or may be divided into other modules Modules and/or sub-modules.

[0174]

In some embodiments, the processing unit 110 can execute instructions stored on one or more modules in the memory 140 or 150, and the execution of the processing unit 110 can cause the system 100 to perform the methods discussed herein.

[0175]

IPR2025-00943
Tesla EX1002 Page 1071

In some embodiments, the processing unit 110 may enable the system 100 to obtain (for example, with the image acquisition module 1100) multiple images from the camera (for example, the image capture device 122, 124, and/or 126), wherein, when the camera is installed in the vehicle (E.g., vehicle 200) and when the vehicle is in motion, the camera is adapted to capture an image of the road within the camera's field of view, and apply image processing to multiple images to identify (e.g., use the road lane restriction recognition module 1102).) The lane restriction of the first road is determined based on multiple images (for example, using the determining module 1104) that one or more defined lane prediction conditions are satisfied, and based on the determination that one or more defined lane prediction conditions are satisfied, predict (for example, Using the road lane restriction prediction module 1106) the position of the second road lane restriction, and issuing (for example, using the issuing module 1108) notification or control signal according to the predicted position of the second road lane restriction to control the movement of the vehicle.

[0176]

In some embodiments, the processing unit 110 may cause the system 100 to issue (for example, using the issuing module 1108) notification or control signal according to the location of the lane restriction on the first road to guide the movement of the vehicle.

[0177]

In some embodiments, the processing unit 110 may cause the system 100 to issue (for example, using the issuing module 1108) notification or control signal to guide the vehicle along the lane or area between the lane restriction on the first road and the lane restriction on the second road. sports.

[0178]

In some embodiments, the processing unit 110 may cause the system 100 to determine (for example, using the determining module 1104) whether one or more defined lane prediction cancellation conditions are met, and predict whether the cancellation conditions are met based on the one or more defined lane cancellation conditions. Determine, stop predicting (for example, using the lane restriction prediction stop module 1110) the position of the lane restriction on the second road.

[0179]

Figure 12 illustrates a process 1200 for implementing lane restriction prediction consistent with the disclosed embodiment.

In some embodiments, the steps in the process 1200 may be illustrated by the situations illustrated in FIGS. 8A-8C, 9A-9C, and 10A-10B.

According to some embodiments, the process 1200 may be implemented by one or more components of the navigation system 100 (such as at least one processing unit 110).

According to some embodiments, the process 1200 may be implemented by the system 100 associated with the vehicle 200.

[0180]

In step 1202, in some embodiments, the system receives multiple images of roads near the moving vehicle associated with the system.

In some embodiments, one or more images are captured using at least one image capture device (such as image capture devices 122, 124, and/or 126).

Multiple images may be received by the system through a data interface such as data interface 128.

In the example illustrated in FIGS. 8A-8C, the system 100 may collect multiple images of the road 800 near the vehicle 200.

[0181]

In step 1204, in some embodiments, the system identifies the first road lane restriction.

Tesla EX1002 Page 1072 IPR2025-00943

In some embodiments, the system can use image analysis techniques to detect objects, markers, or other characteristics (such as lane markings, road edges, lane drivers, and/or other vehicles) in images that indicate lane restrictions. One or more images received by lane restrictions are processed.

In some embodiments, the system may attempt to identify at least one lane restriction on one side of the vehicle and another lane restriction on the other side of the vehicle (e.g., in the front-to-rear direction of the vehicle, its centerline One lane limit on each side). In some embodiments, the system may attempt to identify at least one lane restriction on one side of the lane and another lane restriction on the other side of the lane. In the example illustrated in FIGS. 8A-8C, the system 100 may analyze the received image of the road 800, and may detect the lane restriction 804 on the left side of the lane 802 corresponding to and matching the real lane line 803a. That is, in some embodiments, the lane line 803a can be recognized by the system 100 in a plurality of captured images, and the system 100 can determine that the lane line 803a is a visual mark for the lane restriction 804.

[0182]

In step 1206, in some embodiments, the system determines whether one or more defined lane prediction conditions are met based on the multiple images.

In some embodiments, the lane prediction condition may be any predetermined condition, and the system 100 can determine whether to predict the location of the lane restriction through the lane prediction condition. In some embodiments, only one lane prediction condition needs to be met, while in other embodiments, the minimum threshold number of lane prediction conditions (e.g., two or more) needs to be met. In some embodiments, the minimum threshold number of lane prediction conditions may be based on the relationship between the road lane layout, road topology, shadow conditions, visibility conditions, the presence of guardrails and/or other obstacles along the road, and other vehicles on the road. Existence and other relevant historical information to determine.

[0183]

In some embodiments, the one or more defined lane prediction conditions include a second road lane restriction stop being identified based on multiple images as the vehicle progresses.

For example, in some embodiments, the system can continuously capture and refer to multiple images corresponding to the current location and vicinity of the vehicle; as the vehicle travels along the road and along the lane, a subset of the multiple images can be In this case, the lane restriction that can be recognized in time can be based on the subsequent subset of multiple images to stop being recognized in time in the subsequent case. In the example of FIGS. 8A-8C, as the vehicle 200 travels along the road 800 from its location in FIG. 8A to its location in FIG. 8B, the system 100 may analyze the received image of the road 800. As the vehicle 200 travels from its location in FIG. 8A to its location in FIG. 8B, the vehicle 200 reaches the middle portion of the road 800 where the lane line 803 b dotted in the middle portion is not drawn on the road 800. When analyzing the multiple images received from the middle part of the road 800, the system 100 may stop recognizing the lane restriction 806 due to the absence of the lane line 803b. This stop-recognition lane restriction as the vehicle 200 advances along the road 800 may constitute a lane prediction condition in some embodiments.

[0184]

In some embodiments, the one or more defined lane prediction conditions include that the lane restriction on the first side of the lane has been recognized for at least a threshold time or a threshold distance.

Optionally, the threshold time or threshold distance associated with one or more defined lane prediction conditions can be dynamic, and can be based on road lane layout, road topology, shadow conditions, visibility conditions, guardrails along the road Or the existence of other obstacles, the existence of other vehicles on the road and other relevant historical information to adjust. In some embodiments, such lane prediction conditions can be used to prevent the system from performing lane restriction predictions (or reduce the system's lane restrictions when the system detects lane restrictions on a single side of the lane (or a single side of the vehicle) for the first time. Limit the possibility of prediction). That is, in some embodiments, if the vehicle is traveling and encounters a single lane line or other lane markings or indicators for the first time, the system may not predict the location of another lane restriction based on the single lane marking/indicator . In some embodiments, the lane prediction conditions can be used to ensure that the system makes effective and accurate lane restriction predictions based on the actual lanes present on the road rather than isolated, stray, or incorrect or randomly placed lane markings/indicators (Or increase its possibility). In some embodiments, the predetermined threshold distance may be fixed in the system, or may be

passed based on the current speed of the vehicle, the size of the road on which the vehicle travels or the marked speed, or any other discernible variable related to the vehicle or road. The system is determined.

[0185]

In some embodiments, the one or more defined lane prediction conditions include the presence of a guardrail near the vehicle.

In some embodiments, when the system determines that the vehicle is near the guardrail, the system can make a lane restriction prediction. In some embodiments, the lane prediction condition can be used to prevent the system from erroneously determining one or more detected in one or more images when one or more edges are associated with a guardrail or with a guardrail shadow in reality. The edges correspond to lane restrictions. In some embodiments, the guardrail can cast shadows on the road, and the edges of the shadows on the road can be detected by the image analysis technology used by the system. In some embodiments, the system may erroneously determine that the guardrail shadow is a sign of lane restrictions, and may cause the vehicle to travel based on the guardrail shadow instead of the actual lane restrictions. In some embodiments, in order to prevent such errors, the system may recognize that the guardrail is present in the vicinity of the vehicle, and thus may perform lane restriction prediction to predict the position of the lane restriction on the side of the lane with the guardrail, instead of trying to identify the lane Of the lane restriction on that side and risked wrongly identifying the shadow of the guardrail as a lane restriction.

[0186]

In the example of FIGS. 10A and 10B, (shown in FIG. 10B; not shown in FIG. 10A, because the vehicle 200 constitutes a viewpoint in FIG. 10A) the vehicle 200 travels in the lane 1002 on the road 1000, the The lane has a guardrail 1012 on the left side of lane 1002.

In some embodiments, the system 100 may analyze one or more images to detect that the guardrail 1012 is present in the vicinity of the vehicle 200 on the left side of the lane 1002. In some embodiments, the presence of the guardrail 1012 can satisfy the lane prediction conditions.

[0187]

In some embodiments, determining that the guardrail is present in the vicinity of the vehicle includes detecting in one or more images a pattern created by the pillar supporting the guardrail.

In some embodiments, the system may be configured to detect one or more patterns in one or more images that are known or suspected to indicate the presence of a guardrail. In some embodiments, one such pattern may be referred to as a "rapid dash" pattern, and when the system detects in one or more images a rapid continuous dash or rapid continuous dash generated by the pillar supporting the guardrail When the line is used, such a pattern may exist. In some embodiments, the system may be configured to detect any other suitable patterns or visual markings known to be characteristics of the pillar supporting the guardrail. In some embodiments, the system can detect the pattern of the posts of the guardrail, and in some embodiments, alternatively or additionally, the system can detect the pattern of shadows cast by the posts of the guardrail. In the example of FIGS. 10A and 10B, the guardrail 1012 is supported by a column 1014 (shown in FIG. 10A; not visible from the overhead angle of FIG. 10B). In FIGS. 10A and 10B, the pillar 1014 casts a shadow that falls on the road 1000 and in the lane 1002. The system 100 can detect the presence of a quick dotted line pattern or any other suitable pattern created by the pillar 1014 or its shadow in multiple images, and this detection can be a form in which the system 100 determines that a guardrail is present in the vicinity of the vehicle 200.

[0188]

In some embodiments, the system 100 can use the depth information to recognize the shape of the guardrail or the characteristic of the part of the guardrail.

It should be appreciated that the drawn line will have substantially no volume and its height relative to the road surface is almost zero, so depth information can be used to distinguish between lane markings and guardrails.

[0189]

In some embodiments, the one or more defined lane prediction conditions include guardrails casting shadows onto the road.

IPR2025-00943
Tesla EX1002 Page 1074

In some embodiments, instead of or in addition to the detection of the guardrail itself by the system, the lane prediction condition may also be based on the detection of the shadow of the guardrail by the system. In some embodiments, if the guardrail shadow is detected anywhere in the vicinity of the vehicle, the lane prediction condition may be met; in some embodiments, if the shadow is detected as coming from the lane in or on the vehicle Or the projection of the guardrail in the direction of the road, the lane prediction condition may be satisfied; in some embodiments, if the shadow is detected as being projected from the guardrail onto the road or lane in which the vehicle is traveling, the lane prediction condition may be satisfy.

[0190]

In some embodiments, determining that the guardrail casts a shadow onto the road may involve determining whether the predicted position of the sun is consistent with the position of the detected edge that may be the shadow of the guardrail.

For example, in some embodiments, the system may determine the orientation of the sun relative to the multiple images available to the system and the position of the sun. The determination of the position of the sun and the orientation of the vehicle can be based on various pieces of information (for example, including one or more of multiple images, the position and/or orientation of the vehicle determined by GPS, the orientation of the vehicle determined by the compass And/or time) to complete. The system can determine that the shadow is expected to be cast from the guardrail in a certain direction at a certain time based on the determined position of the sun and the orientation of the vehicle.

[0191]

In the example of FIGS. 10A and 10B, the guardrail 1012 casts a shadow 1016 onto the road 1000 and into the lane 1002; the shadow 1016 that falls on the road 1000 and falls in the lane 1002 can be detected by the system 100, and the The detection can meet the lane prediction conditions.

[0192]

In some embodiments, the one or more defined lane prediction conditions include multiple substantially parallel edges being detected in multiple images on the road on one side of the lane on the road.

In some embodiments, the image analysis techniques applied by the system to the received multiple images can detect edges in one or more of the received multiple images; in some embodiments, edge detection can be used to identify Objects and markings in the environment around the vehicle.

In some embodiments, the detection of edges in multiple images can be used to identify lane lines or markings or edges or roads, and thus to determine the location of lane restrictions. In some embodiments, including those in which the guardrail shadow is projected onto the road or lane on which the vehicle travels, the system can detect the edge of the guardrail shadow as an edge in multiple images. In some other embodiments, other objects or markings may enable the system to detect multiple edges on the same side of a lane. In these embodiments, the determination of which edge corresponds to the correct lane restriction may be difficult, slow, or unreliable. For this reason, in some embodiments, it may be advantageous not to try to determine which of several detected edges corresponds to a lane restriction based on multiple images, and as discussed above, when multiple When edges are detected in multiple images, lane restriction prediction is used instead of simply predicting the location of the lane restriction. In the example of FIGS. 10A and 10B, the system 100 detects three different edges indicated by bold lines 1004, 1005a, and 1005b on the left side of the lane 1002 in multiple images. In the example of FIGS. 10A and 10B, as explained above, the system 100 may determine that multiple edges are detected on the left side of the lane 1002, and this detection may constitute a lane prediction condition.

[0193]

In some embodiments, if it is determined at block 1206 that one or more of the defined lane prediction conditions are not met, the method 1200 may return to step 1202 and may continue to acquire additional images of the road over time.

[0194]

In some embodiments, if it is determined at block 1206 that one or more of the defined lane prediction conditions are satisfied, the method 1200 may proceed to step 1208.

In step 1208, in some embodiments, based on the determination that one or more defined lane prediction conditions are satisfied, the system predicts the location of the lane restriction on the second road.

In some embodiments, the system may predict the lane restriction on the opposite side of the lane from the detected lane restriction (the vehicle serves as a reference for the direction of the opposite side). Therefore, in some embodiments, if a lane restriction can be identified from lane lines or other visual signs in the received image, and it is determined that the lane prediction conditions such as any of the conditions discussed above are met, the system may Predict the position of the second lane restriction on the side of the lane opposite to the first lane restriction.

[0195]

In some embodiments, the predicted lane restriction may be predicted to extend along the road within a predetermined distance.

In some embodiments, the predicted lane restriction may be predicted to be parallel to the identified lane restriction opposite it. In some embodiments, the predicted lane restriction may be predicted to continue along the same angle, path, or location as the same lane restriction previously detected by the system. In some embodiments, the predicted lane restriction may be predicted based on the location of another lane restriction, based on the location of other distinguishable features of the road, or based on the location of vehicles associated with the system and/or surrounding vehicles. In some embodiments, the predicted lane limit may be predicted to be located at a predetermined distance from the identified lane limit or from any other objects identified or predicted based on multiple images of the road. In some embodiments, the predicted lane restriction may be predicted to match or be determined based on other lane restriction types detected in multiple images, or be predicted to be based on roads from a previous time. The previous image of the lane restriction type determined by the lane restriction type determined by the predicted lane restriction (for example, if the system detects a real single white line on the right side of the vehicle, it will stop detecting any lane line, and the system can predict The solid single white lane is restricted to continue on the right side of the vehicle; alternatively or additionally, the system may predict that the solid single white line is located on the side of the lane opposite the solid single yellow line).

[0196]

In the example of FIGS. 8A-8C, the system 100 recognizes lane restrictions 804 based on the detection of lane lines 803a in multiple images.

When the vehicle 200 reaches its location in FIG. 8B and no lane line is detected on the right side of the lane 802, the lane prediction condition may be satisfied. Therefore, in some embodiments, the system 100 may predict the lane restriction 806 on the right side of the lane 802, as shown in FIG. 8B. In some embodiments, the lane restriction 806 may be predicted as a single dashed lane restriction based on the previous recognition in FIG. 8A based on the lane line 803b of the lane restriction 806 that is a single dashed lane restriction. In some embodiments, the location, angle, type, and/or any other characteristics of the lane restriction 806 may be based on any stored or otherwise accessible factors about the vehicle 200, the road 800, or any other surrounding environment accessible through the system 100 Deterministic information to predict.

[0197]

In some embodiments, predicting the location of the lane restriction includes predicting the width of the lane based on the width of the previously determined lane.

In some embodiments, the width of the lane may be predicted based on historical information about the width of the lane or cached information about the width of a specific lane that is the latest (for example, a few seconds before the lane restriction prediction is performed). In some embodiments, when the system detects two limits of a lane and thus can determine the width of the lane and the system then stops detecting or identifying one of the lane limits, the system can predict that the lane continues to be the same width. The system can accordingly predict the location of the predicted lane restriction so that the lane is predicted to maintain a constant or substantially constant width. In the example of FIGS. 8A-8C, when the system 100 predicts the location of the lane restriction 806 in FIG. 8B, the restriction 806 is predicted to be a distance 812 from the lane restriction 804. When the vehicle 200 is in the position shown in FIG. 8A, the distance 812 may be the same distance 812 detected between the lane restriction 806 and the lane restriction 804. In some embodiments, the system 100 may record or cache historical information about the lane width, so that the system can quickly adjust the lane width information to predict the position of the predicted lane restriction.

[0198]

In step 1210, in some embodiments, the system enables the moving vehicle to avoid the first road lane restriction and the second road lane restriction.

Therefore, in some embodiments, the system may cause the vehicle to travel in a lane defined on one side of the detected/identified lane restriction and on the opposite side of the predicted lane restriction. In some embodiments, the system may cause any suitable steering, acceleration, deceleration, or other control of the vehicle to cause the vehicle to travel in the lane defined by the predicted and detected lane restrictions. In the example of FIGS. 8A-8C, when the system 100 detects the lane restriction 806 in FIG. 8B, the system 100 then causes the vehicle 200 to follow the lane between the detected lane restriction 804 and the predicted lane restriction 806 802 travels to travel from its location in FIG. 8B to its location in FIG. 8C.

[0199]

In step 1212, in some embodiments, the system determines whether one or more defined lane prediction cancellation conditions are met.

In some embodiments, the lane prediction cancellation condition may be any predetermined condition, and the system can determine whether to stop predicting the location of the lane restriction through the lane prediction cancellation condition. In some embodiments, once the system begins to predict the location of the lane restriction, the location of the lane restriction can be continued to be predicted until one or more predetermined lane prediction cancellation conditions are met. In some embodiments, only one lane prediction cancellation condition needs to be satisfied in order to stop the lane prediction. In some embodiments, the predetermined number or lane prediction condition is greater than a condition that must be met in order to stop the lane prediction. Alternatively, the confidence score may be assigned to the lane prediction conditions according to a predetermined criterion. Optionally, some lane prediction conditions may have a higher weight than others. Still further by way of example, the overall score may be calculated to determine whether the lane prediction cancellation condition is satisfied, and such calculation may consider the weight and confidence score assigned to the lane prediction condition.

[0200]

In some embodiments, the one or more defined lane prediction cancellation conditions include identifying a second lane restriction based on the first or more images after the stop recognition.

For example, in some embodiments, as the vehicle travels along the road, lane restrictions may stop being recognized (as discussed above), such as when the road is not marked or lane lines are not detectable in multiple images; as above As discussed, lane prediction can be performed accordingly. In some embodiments, thereafter, such as when the vehicle continues to travel along the road, as the system continues to acquire new images, the system may recognize the lane restriction in multiple images again. For example, if the vehicle reaches a part of the road that the lane line is no longer missing, or if the lane line becomes detectable again by the system in other ways, the system may identify lane restrictions based on the updated image of the road. The identification of the second lane restriction-the predicted lane restriction-in some embodiments may constitute a lane prediction cancellation condition that can cause the system to stop predicting the location of the second lane restriction and return to relying on multiple The location of the lane restriction detected/recognized by each image.

[0201]

In the example of FIGS. 8A-8C, although the system 100 can predict the position of the lane restriction 806 when the vehicle 200 is at the position shown in FIG. 8B, as the vehicle 200 advances to the position shown in FIG. 8C, the The imaginary lane line 803 on the upper half is again visible in the image collected by the system 100.

Based on these new images, the system 100 can again detect/recognize the presence of the lane restriction 806 based on the lane line 803b, and this recognition can constitute a lane prediction cancellation condition that triggers the stop of the position prediction of the lane restriction 806.

[0202]

In some embodiments, the one or more defined lane prediction cancellation conditions include the driver moving the vehicle between lanes.

IPR2025-00943
Tesla EX1002 Page 1077

In some embodiments, the system may determine the driver to steer or merge the vehicle into another lane based on multiple images. In some embodiments, the detection may satisfy the lane prediction cancellation condition, and may trigger the cancellation of the lane restriction position prediction.

[0203]

In some embodiments, the one or more defined lane prediction cancellation conditions include the vehicle crossing a curve that exceeds a predetermined sharpness threshold.

In some embodiments, the system has access to a predetermined or dynamically determined sharpness threshold of the curve (e.g., determined based on vehicle speed, marked speed limit, road conditions, weather conditions, environmental brightness, proximity of other vehicles, etc.) sharpness threshold, and it can be determined whether the curve in the road or lane exceeds the sharpness threshold (eg, sharper than the maximum allowable sharpness). If it is determined that the curve that the vehicle is crossing or approaching exceeds the maximum allowable sharpness, the system may determine that the lane prediction cancellation condition is satisfied, and may stop the lane restriction prediction accordingly.

[0204]

In some embodiments, determining whether one or more defined lane prediction cancellation conditions are met includes selecting a maximum lane prediction distance setting, and determining whether the vehicle travels a distance greater than the maximum lane prediction when predicting the location of the second road lane restriction The maximum lane prediction distance associated with the distance setting.

In some embodiments, the maximum lane prediction distance may be predetermined, set or dynamically determined by the system. In some embodiments, the system may monitor the distance traveled by the vehicle when the system predicts the location of the lane restriction, and if the travel distance exceeds the maximum allowable predicted lane distance, it may stop predicting the location of the lane restriction. In some embodiments, the implementation of the maximum allowable lane prediction distance can improve safety and functionality by ensuring that the system does not predict the location of lane restrictions during an excessively long period of time or within an excessively long distance, during which distance or time, One or more characteristics of the lane or road may change or change. In some embodiments, since the characteristics of lanes and roads can be expected to change or change with space and time, the system can ensure that the maximum lane prediction distance is exceeded (or, in some embodiments, the maximum lane prediction time) at least once the maximum lane prediction distance is exceeded.) Lane restriction reference points or indicators (such as lane lines or other lane markings) are detected based on multiple images collected by the system to improve safety.

[0205]

In some embodiments, the maximum lane prediction distance can be changed by the system in different situations.

In some embodiments, the maximum lane prediction distance may be associated with one or more maximum lane prediction distance settings, so that the system can select a maximum lane prediction distance setting and implement the associated maximum lane prediction distance. In some embodiments, different maximum lane prediction distance settings may be implemented in different driving scenarios detected and determined by the system.

[0206]

In some embodiments, the standard or default maximum lane prediction distance setting may be associated with the standard or default maximum lane prediction distance.

In some embodiments, the standard or default maximum lane prediction distance may be 18 meters.

[0207]

In the example of FIGS. 8A-8C, the vehicle 200 travels from the driving distance 814 predicted by the lane restriction to its location in FIG. 8C when the lane is predicted to stop.

The distance 814 may be less than the standard or default maximum lane prediction distance of 18 meters, so that the lane restriction prediction may be continuously performed between the positioning of the vehicle 200 in FIGS. 8B and 8C.

IPR2025-00943
Tesla EX1002 Page 1078

[0208]

In some embodiments, selecting the maximum lane prediction distance setting includes, according to the determination that the first lane and the second lane are merged in the vicinity of the vehicle, selecting the maximum lane prediction distance setting so that the maximum lane prediction distance is greater than the default maximum lane prediction distance .

Therefore, in some embodiments, the system may implement an extended or increased maximum lane prediction distance. For example, in some embodiments, the system may determine, based on multiple images, that the lane merges with the lane the vehicle travels, or the lane the vehicle travels merges with another lane. The system can detect lane restrictions on the far side of the merged lane of the lanes of adjacent vehicles (for example, by identifying lane lines) and by determining that the lanes on the opposite side of the adjacent lanes are restricted above a predetermined or dynamically determined approach threshold rate. Approach the lane where the vehicle is driving to determine that the lane merging occurs in the vicinity of the vehicle. That is, in some embodiments, as the vehicle advances along the road, the system may identify the far side of the adjacent lane of the lane approaching the vehicle traveling, and may accordingly determine that the adjacent lane merges with the vehicle traveling lane.

[0209]

In some embodiments, when two lanes merge with each other, it is expected that the road may be characterized by a long-distance unmarked lane restriction, therefore, it is configured to be at a distance greater than the default maximum lane predicted distance or time. And/or systems that predict the location of lane restrictions over a longer period of time may be beneficial.

In some embodiments, the increased maximum lane prediction distance associated with lane merging conditions may be 1000 meters.

[0210]

In the example of FIGS. 9A to 9C, the vehicle 200 has traveled a distance 914 from the location of the lane restriction prediction in FIG. 9B to the location of the lane prediction stop in FIG. 9C.

The distance 914 may be greater than the standard or default maximum lane prediction distance, so that if the default maximum lane prediction distance setting is selected, the lane restriction prediction may have been stopped before the vehicle 200 reaches its location in FIG. 9C. However, as indicated by the bold line in FIG. 8B, as the vehicle 200 advances, the system 100 recognizes the lane restriction 922 on the right side of the lane 915 approaching the lane 902, and can determine accordingly that the lanes are merging, and should choose to The maximum lane prediction distance setting associated with lane merging. The selected setting may be associated with a maximum lane prediction distance of 1000 meters that is greater than the default maximum lane prediction distance of 18 meters. Therefore, in some embodiments, the distance 914 may be less than an increased maximum lane prediction distance of 1000 meters, so that the lane restriction prediction may be continuously performed between the positioning of the vehicle 200 in FIG. 9B and FIG. 9C.

[0211]

In some embodiments, selecting the maximum lane prediction distance setting includes, according to a determination that the defined guardrail condition is satisfied, selecting the maximum lane prediction distance setting without the maximum lane prediction distance.

That is, in some embodiments, when the system 100 detects that the guardrail condition is satisfied, the system 100 can predict the location of the lane restriction indefinitely without considering any maximum lane prediction distance. In this way, in some embodiments, the maximum lane prediction distance associated with the guardrail condition may be considered an infinite maximum distance or a non-existent maximum distance.

[0212]

In some embodiments, the guardrail condition may be satisfied in any situation where the system detects the presence of one or more guardrails, guardrail posts, and/or guardrail shadows in the vicinity of the vehicle.

For example, the guardrail conditions can be met in any or all of the above-discussed situations, where the system can determine that the guardrail and/or guardrail shadow exists in the vicinity of the vehicle and/or on the road or lane on which the vehicle travels. As discussed above, in some cases, the guardrail condition may be satisfied when the presence, location, and/or orientation of the guardrail and/or guardrail shadow cause the system to make lane restriction predictions.

[0213]

In some embodiments, it may be advantageous not to apply the maximum lane prediction distance (or infinite maximum lane prediction distance) to the situation where the guardrail and/or guardrail shadow exists, because the guardrail and its associated shadow are expected to follow along The road can continue for a very long distance (typically greater than the default maximum lane prediction distance or even greater than the increased lane prediction distance associated with lane merging).

In the example of FIGS. 10A and 10B, when the detection system 100 starts to predict the location of the lane restriction 1004 due to the guardrail shadow 1016, the system 100 may select the maximum lane prediction distance setting associated with the guardrail condition; The selection of the maximum lane prediction distance setting enables the system to perform lane prediction of lane limit 1004 without any maximum lane prediction distance as the premise.

[0214]

In some embodiments, if it is determined at block 1212 that one or more of the defined lane prediction cancellation conditions are not met, the method 1200 may return to step 1202 and may continue to acquire additional images of the road over time.

[0215]

In some embodiments, if it is determined at block 1212 that one or more of the defined lane prediction cancellation conditions are met, the method 1200 may proceed to step 1214.

In step 1212, in some embodiments, based on the determination that one or more lane prediction cancellation conditions are satisfied, the system stops predicting the location of the second lane restriction.

In the example of FIGS. 8A-8C, when the vehicle 200 reaches its position depicted in FIG. 8C, the lane prediction cancellation condition may be satisfied based on the recognition of the lane line 803b passing the lane restriction 806, and therefore the system 100 may stop the prediction The location of lane restriction 806. Conversely, as the system 100 did previously in FIG. 8A, it can re-determine the location of the lane restriction 806 based on the detected location of the lane line 803b.

[0216]

Detection and response in construction area

[0217]

The system 100 may provide construction area detection to the vehicle 200, so that the system recognizes the existence of the construction area on the road and takes actions on it.

In some embodiments, the system 100 can identify various elements in one or more images it receives, and can determine that the vehicle 200 is approaching or in the construction area based on the identification of these one or more elements.

In some embodiments, any automatic or driver-assisted actions can be taken based on the detection of the construction area, including but not limited to changing the speed of the vehicle 200, turning the vehicle 200, activating or disabling one or more lights of the vehicle 200, and so on. In some embodiments, based on the detection of the construction area, the signal may be output by the system 100, which indicates that the construction area has been detected and/or the construction area mode is activated. In some embodiments, according to the detection of the construction area (and/or according to the detection of the complex construction area), the automatic control of the vehicle 200 or the control of the driver assistance performed by the system 100 may be disabled, so that the control of the vehicle 200 Can be returned to the driver.

[0218]

FIG. 13 illustrates an embodiment in which the vehicle 200 travels on the road 1300 and approaches the construction area.

IPR2025-00943
Tesla EX1002 Page 1080

The road 1300 has two lanes 1302 and 1304. Lane 1302 is defined by lane restrictions 1306 and 1308, and lane 1304 is defined by lane restrictions 1310 and 1312. As indicated by the bold lines indicating lane restrictions 1306, 1308, 1310, and 1312, these lane restrictions are detected by the system 100. The lane restrictions 1306 and 1308 are marked on the road 1300 through the drawn lane lines 1320 and 1322, respectively, and the lane restrictions 1310 and 1312 are marked on the road 1300 through the drawn lane lines 1322 and 1324, respectively.

[0219]

The road 1300 includes several indicators of the construction area or construction area, including the drawn construction area lane line 1326.

In some embodiments, the lane line of the construction area may be drawn in a different color and/or style from the traditional lane line. For example, in Germany, yellow lane lines are drawn on the road to indicate construction areas, while traditional lane lines are drawn in white. As indicated by the bold line, the system 100 detects the lane restriction 1314 of the construction area based on the drawn lane line 1326 of the construction area. According to the detection of the lane restriction 1314 in the construction area, the system 100 can enable the vehicle 200 to drive in the lane defined by the lane restriction 1314 in the construction area instead of the traditional lane restriction. In some embodiments, when the system 100 detects that the vehicle 200 is in the construction area, the construction area lane line may take precedence over the traditional lane line.

[0220]

The road 1300 may include traditional construction area indicators, including cones 1316, some of which are located on the surface of the road 1300 and within the lane 1304.

In response to detecting the cone 1316, the system 100 may cause the vehicle 200 to change its path in order to avoid the cone 1316, change the position in the lane of travel, or change the lane if necessary. In response to the detection cone 1316, the system 100 may determine that the vehicle 200 is in or near the construction area; activate the construction area mode; output the construction area signal; disable the automatic control or driver assistance control of the vehicle 200; and/or by changing it Speed, turn, or take any other appropriate action (such as avoiding the cone 1316) to control the vehicle 200.

[0221]

The road 1300 may include additional construction zone indicators, including construction road signs 1318a and 1318b, which can indicate the existence of a construction zone and/or instructions for the driver to follow in the construction zone (e.g., slow down, stop, merge, Turn, etc.).

The system 100 can detect the construction road signs 1318a and 1318b in a manner in which the system can detect any or all other traffic signs, and the system can take any appropriate action in response to the detection of the construction road signs 1318a and 1318b, including determining that the vehicle 200 is in or near the construction area; activating the construction area mode; outputting the construction area signal; disabling the automatic control or driver assistance control of the vehicle 200; and/or by changing its speed, steering or taking any other appropriate Actions (such as in accordance with instructions on construction road signs) control the vehicle 200.

[0222]

FIG. 14 is an exemplary block diagram of the memory 140 or 150, which may store instructions for performing one or more operations consistent with the disclosed embodiments.

As illustrated in FIG. 14, the memory 140 or 150 may store one or more modules for performing the construction area detection described herein. For example, the memory 140 or 150 may store an image acquisition module 1400, a construction area indicator recognition module 1402, a determination module 1404, a signal output module 1406, a vehicle movement enabling module 1408, and a vehicle control disabling module 1410. As those of ordinary skill in the art will recognize, any of the modules listed above can be combined with each other module or other modules, joined with another module or other modules, and/or divided into additional modules And/or submodules.

[0223]

IPR2025-00943
Tesla EX1002 Page 1081

In some embodiments, the processing unit 110 can execute instructions stored on one or more modules in the memory 140 or 150, and the execution of the processing unit 110 can cause the system 100 to perform the methods discussed herein.

[0224]

In some embodiments, the processing unit 110 may enable the system 100 to obtain (for example, with the image acquisition module 1400) multiple images from a camera (for example, the image capture device 122, 124, and/or 126), wherein, when the camera is installed in the vehicle (For example, the vehicle 200) and when the vehicle is in motion, the camera is adapted to capture images of roads in its field of view, and apply image processing to multiple images to identify (for example, using the construction area indicator recognition module 1402) One or more indicators in the construction area, and according to the one or more indicators identifying the construction area, output (for example, using the signal output module 1406) to output a signal indicating that the vehicle is approaching the construction area.

[0225]

In some embodiments, the processing unit 110 may cause the system 100 to determine (for example, using the determination module 1404) that the vehicle is approaching the construction area based on one or more indicators, and output (for example, using the signal output Module 1406) A signal indicating that the vehicle is approaching the construction area.

[0226]

In some embodiments, the processing unit 110 may cause the system 100 to move the moving vehicle (for example, using the vehicle movement activation module 1408) according to the construction area indicator according to the determination that the vehicle is approaching the construction area.

[0227]

In some embodiments, the processing unit 110 may enable the system 100 to determine (for example, using the determination module 1404) that the construction area is a complex construction area based on one or more indicators, and according to the determination that the construction area is a complex construction area, disable (For example, using the vehicle control disable module) automatic vehicle control functions.

[0228]

Figure 15 illustrates a process 1500 for implementing lane restriction prediction consistent with the disclosed embodiment.

In some embodiments, the steps of the process 1500 may be illustrated by the situation illustrated in FIG. 13.

According to some embodiments, the process 1500 may be implemented by one or more components of the navigation system 100 (such as at least one processing unit 110).

According to some embodiments, the process 1500 may be implemented by the system 100 associated with the vehicle 200.

[0229]

In step 1502, in some embodiments, the system receives multiple images of roads near the moving vehicle associated with the system.

In some embodiments, one or more images are captured using at least one image capture device (such as image capture devices 122, 124, and/or 126).

Multiple images may be received by the system through a data interface such as data interface 128.

In the example illustrated in FIG. 13, the system 100 may collect multiple images of a road 1300 near the vehicle 200.

[0230]

In step 1504, in some embodiments, the system identifies one or more indicators of the construction area in multiple images.

In some embodiments, the system may use image analysis techniques to process one or more of the received images to detect signs, objects, other vehicles, marked lane lines, and any other objects or characteristics of the surrounding environment.

In some embodiments, certain signs, markings, lane lines, objects, and/or combinations of any such detected aspects may be predetermined by the system as indicators of construction zones. In some embodiments, cones, cylinders, flashing lights, poles, obstacles, and/or markings present near the road can be predetermined by the system as indicators of the construction zone. In some embodiments, predetermined traffic signs, such as those carrying words, symbols, or images known to indicate the construction area, can be predetermined by the system as indicators of the construction area. In some embodiments, the lane line may be predetermined by the system as an indicator of the construction area. For example, in Germany, the construction area can be marked with a special construction area lane line drawn on the road. The construction area lane line may be a different color or a different appearance from the traditional lane line. The system can detect the color, appearance, orientation, and/or position of the lane line, which indicates to the system that the line is a lane line in the construction area, and the system can recognize such a lane line in the construction area as an indicator of the construction area.

[0231]

In the example of FIG. 13, the system 100 can identify the cone 1316, the construction road signs 1318a and 1318b, and/or the drawn construction area lane line 1326.

In some embodiments, any one or all of these elements may be predetermined by the system 100 or dynamically determined by the system 100 as a construction area indicator.

[0232]

In step 1506, in some embodiments, the system determines whether the vehicle is approaching the construction area based on one or more indicators.

In some embodiments, the system may determine that the vehicle is approaching the construction area based on the vicinity (of the vehicle, road, or lane) of the construction area indicator. In some embodiments, the system may assume that the vehicle is approaching the construction area based on the orientation or location relative to the road or lane. In some embodiments, the system may determine that the vehicle is approaching the construction area based on the content of one or more road signs (for example, images, symbols, and/or text). In some embodiments, the minimum threshold number of construction zone indicators must be detected within the maximum threshold time of each other in order for the system to determine that the vehicle is approaching the construction zone. In the example of FIG. 13, the system 100 may simultaneously detect and determine that the vehicle 200 is approaching the construction area based on the cone 1316, the construction road signs 1318a and 1318b, and/or the drawn construction area lane line 1326.

[0233]

In some embodiments, if it is determined at block 1506 that the vehicle is not approaching the construction area, the method 1500 may return to step 1502 and may continue to collect additional images of the road over time.

[0234]

In some embodiments, if it is determined at block 1506 that the vehicle is approaching the construction area, the method 1500 may proceed to step 1508.

In step 1508, in some embodiments, based on the determination that the vehicle is approaching the construction area, the system outputs a signal indicating that the vehicle is approaching the construction area.

In some embodiments, the signal output by the system 100 may be an electrical signal or an electrical message, which informs any component of the system or any associated system that the construction area is detected. For example, the signal can be sent to a display system associated with the system, or the signal can be sent to any other local system or remote system that is communicatively coupled with the system. In some embodiments, the signal may be perceivable by the driver of the vehicle. For example, the signal may be a visible signal (such as a light-emitting indicator or display element on a computer or entertainment system display), an audible signal emitted by a speaker system, and/or a tactile signal such as vibration.

[0235]

In step 1510, in some embodiments, based on the determination that the vehicle is approaching the construction area, the system enables the moving vehicle to move according to the construction area indicator.

In some embodiments, in response to the determination that the vehicle is approaching the construction area, the system may cause the vehicle to accelerate, decelerate, turn, or otherwise control and move according to the detected construction area indicator. In some embodiments, the vehicle may automatically decelerate to a reduced speed to increase safety. In some embodiments, vehicles can be made to turn or change lanes in order to follow the lane line of the construction area, avoid cones/cylinders/rods/obstacles or other construction area indicators, or follow or be associated with the indicated construction area The instructions indicated on the road signs. In the example of FIG. 13, the system 100 may modify the progress of the vehicle 200 so that the vehicle 200 follows the detected lane restriction 1314 according to the lane line 1326 of the construction area (rather than the detected lane restriction 1312 according to the lane line 1324) , Avoid the cone 1316 and obey the traffic signs 1318 (which include a symbol that instructs the vehicle 200 to merge to the left). In some embodiments, the vehicle 200 may be made to merge from the lane 1304 to the left into the lane 1302 accordingly.

[0236]

In step 1512, in some embodiments, the system determines whether the construction area is a complex construction area based on one or more indicators.

In some embodiments, a complex construction area may be a construction area with one or more predetermined or dynamically determined characteristics. The predetermined or dynamically determined characteristics are associated with the construction area through the system, and vehicles in the construction area The autonomous control and/or driver assistance control of the vehicle may be unsafe, where the vehicle's autonomous control and/or driver assistance control may be unreliable or unreliable, or where the vehicle's control should be returned to the vehicle's driver. In some embodiments, a complex construction area may be a construction area where vehicles are driving on or in the lane or on the road. In some embodiments, the complex construction area may be a predetermined or dynamically determined threshold number of cones/cylinders/rods/obstacles or other construction area indicators that are detected within a threshold distance or threshold time from each other. Construction area. That is, in some embodiments, when the construction area is densely filled with construction area indicators, it can be determined as a complex construction area. In the example of FIG. 13, in some embodiments, the road 1300 may be based on the presence of cones 1316 on the road 1300, the cones 1316 in the lane 1304, the minimum threshold number of cones 1316 detected, and the detected The minimum threshold number of total construction area indicators reached is determined by the system 100 to include complex construction areas.

[0237]

In some embodiments, if it is determined at block 1512 that the construction area is not a complex construction area, the method 1500 may return to step 1502 and may continue to collect additional images of the road over time.

[0238]

In some embodiments, if the construction area is determined to be a complex construction area at block 1512, the method 1500 may proceed to step 1514.

In step 1514, in some embodiments, based on the determination that the construction area is a complex construction area, the system may disable the automatic vehicle control function.

In some embodiments, driver assistance functions may be disabled. In some embodiments, the fully automated driving function may be disabled. In some embodiments, all control of the vehicle 200 (eg, acceleration, deceleration, steering, etc.) may be returned to the driver of the vehicle 200 by the system 100.

[0239]

The above description is presented for illustrative purposes.

The description is not exhaustive, but is not limited to the precise form or embodiment disclosed. From the specifications and practical considerations of the disclosed embodiments, modifications and adjustments will be obvious to those skilled in the art. In addition, although aspects of the disclosed embodiments are expressed as being stored in a memory, those skilled in the art will recognize that these aspects may also be stored in other types of computer-readable media such as secondary storage

devices, such as , Hard disk or CD ROM, or other forms of RAM or ROM, USB media, DVD, Blu-ray, or other optical drive media.

[0240]

Computer programming based on written descriptions and published methods is within the skills of experienced developers.

Various programs or program modules can be created using any technology known to those skilled in the art, or can be designed in combination with existing software. For example, program segments or program modules can be adopted or passed. Net Framework,. NetCompact Framework (and related languages, such as Visual Basic, C, etc.), Java, C++, object-oriented C, HTML, a combination of HTML/AJAX, XML, or HTML including Java applets. In addition, it should be understood that the functions of the program modules described in this article are not limited to the specific types of program modules named in this article, but can be implemented in any suitable manner among one or more modules or other software structures. Programming.

[0241]

In addition, although illustrative embodiments have been described herein, the scope of any one or all of the embodiments has as those skilled in the art will recognize based on this disclosure. Effective elements, modifications, omissions, combinations (e.g., combinations of aspects across various embodiments), adjustments, and/or changes.

The limitations in the claims are widely interpreted based on the language used in the claims, and are not limited to the examples described in this specification or during the implementation of the application. The examples are interpreted as non-exclusive. In addition, the steps of the disclosed method can be modified in any way, including by reordering steps and/or inserting or deleting steps. Therefore, it is intended that the specification and examples are considered for illustration only, and the true scope and spirit are indicated by the full scope of the following claims and their equivalents.

IPR2025-00943
Tesla EX1002 Page 1085

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

CLAIMS CN107111742A

1.

A device that can be installed on a vehicle, the device comprising:

A camera, when the camera is installed on the vehicle and the vehicle is in motion, it is suitable for capturing multiple images of the road in the field of view of the camera;

The processor, which is configured as:

Obtaining the plurality of images from the camera;

Applying image processing to the plurality of images to identify a first road lane restriction appearing in at least one of the plurality of images;

Based on the multiple images, determining that one or more defined lane prediction conditions are satisfied;

According to the determination that one or more defined lane prediction conditions are met, predict the location of the lane restriction on the second road;

According to the predicted position of the lane restriction on the second road, a notification or a control signal is issued to control the movement of the vehicle.

2.

The device according to claim 1, wherein the processor is configured to issue a notification or a control signal to guide the movement of the vehicle according to the position of the lane restriction on the first road.

3.

The device of claim 1, wherein the processor is configured to issue a notification or a control signal to guide the vehicle along the lane between the first road lane restriction and the second road lane restriction Or regional sports.

4.

The device according to claim 1, wherein the processor is configured to issue a notification or control signal for moving the vehicle, and issue a notification or control signal for guiding the movement of the vehicle to avoid The lane restriction on the first road and the lane restriction on the second road.

5.

The device according to claim 1, wherein the one or more defined lane prediction conditions include, as the vehicle moves on the road, based on the plurality of images, the second road lane restriction stops identified.

6.

The device of claim 1, wherein the one or more defined lane prediction conditions include that the first road lane restriction is recognized for at least a threshold time or a threshold distance.

7.

The apparatus of claim 1, wherein the one or more defined lane prediction conditions include the presence of a guardrail in the vicinity of the vehicle.

8.

The apparatus of claim 7, wherein determining that a guardrail is present in the vicinity of the vehicle includes detecting a pattern created by a pillar supporting the guardrail in the one or more images.

9.

The apparatus of claim 7, wherein the one or more defined lane prediction conditions include the guardrail casting a shadow onto the road.

10.

The apparatus of claim 1, wherein the one or more defined lane prediction conditions include a plurality of substantially parallel lanes detected in the plurality of images on one side of the lane on the road edge.

11.

The apparatus of claim 1, wherein predicting the position of the lane restriction on the second road includes considering the position of the lane restriction on the first road.

12.

The apparatus of claim 1, wherein predicting the position of the lane restriction of the second road includes predicting the width of the lane based on information about the previously determined width of the lane.

13.

The device of claim 1, wherein the processor is configured to:

Determine whether one or more of the defined lane prediction cancellation conditions are met, and

According to the determination that one or more defined lane prediction cancellation conditions are satisfied, stop predicting the position of the lane restriction on the second road.

14.

The apparatus of claim 13, wherein the one or more defined lane prediction cancellation conditions include, after a stop is recognized, the second road lane restriction is recognized based on the one or more images.

15.

The apparatus of claim 13, wherein the one or more defined lane prediction cancellation conditions include a driver moving the vehicle between lanes.

16.

IPR2025-00943
Tesla EX1002 Page 1087

The apparatus of claim 13, wherein the one or more defined lane prediction cancellation conditions include the vehicle traversing a curve that exceeds a predetermined sharpness threshold.

17.

The apparatus of claim 13, wherein determining whether one or more defined lane prediction cancellation conditions are satisfied comprises:

Select the maximum lane prediction distance setting, and

It is determined whether the distance traveled by the vehicle when predicting the position of the lane restriction on the second road is greater than the maximum lane prediction distance associated with the maximum lane prediction distance setting.

18.

The device according to claim 17, wherein selecting the maximum lane prediction distance setting comprises, according to the determination that the first lane and the second lane are merged in the vicinity of the vehicle, selecting the maximum lane prediction distance setting so that the maximum lane prediction The distance is greater than the default maximum lane prediction distance.

19.

The device according to claim 17, wherein selecting the maximum lane prediction distance setting includes selecting the maximum lane prediction distance setting without the maximum lane prediction distance based on the determination that the defined guardrail condition is satisfied.

20.

The device according to claim 17, wherein selecting the maximum lane prediction distance setting comprises, according to the determination that the defined guardrail condition is satisfied, selecting the maximum lane prediction distance setting so that the maximum lane prediction distance is greater than the default maximum lane prediction distance .

21.

A computer system including:

Memory for storing instructions, and

A processor, which executes the instructions to cause the system to:

Obtaining a plurality of images from a camera, the camera being adapted to capture an image of a road in the field of view of the camera when the camera is mounted on the vehicle and the vehicle is in motion;

Applying image processing to the plurality of images to identify a first road lane restriction appearing in at least one of the plurality of images;

Based on the multiple images, determining that one or more defined lane prediction conditions are satisfied;

According to the determination that one or more defined lane prediction conditions are met, predict the location of the lane restriction on the second road;

According to the predicted position of the lane restriction on the second road, a notification or a control signal is issued to guide the movement of the vehicle.

22.

One method includes:

At a system including a memory storing instructions and a processor executing the instructions:

Obtaining a plurality of images from a camera, the camera being adapted to capture an image of a road in the field of view of the camera when the camera is mounted on a vehicle and the vehicle is in motion;

Applying image processing to the plurality of images to identify a first road lane restriction appearing in at least one of the plurality of images;

Based on the multiple images, determining that one or more defined lane prediction conditions are satisfied;

According to the determination that one or more defined lane prediction conditions are met, predict the location of the lane restriction on the second road;

According to the predicted position of the lane restriction on the second road, a notification or a control signal is issued to guide the movement of the vehicle.

23.

A vehicle including:

main body;

A camera adapted to capture multiple images of a road in the field of view of the camera when the camera is installed on the vehicle and the vehicle is in motion;

The processor, which is configured as:

Acquiring multiple images from the camera;

Applying image processing to the plurality of images to identify a first road lane restriction appearing in at least one of the plurality of images;

Based on the multiple images, determining that one or more defined lane prediction conditions are satisfied;

According to the determination that one or more defined lane prediction conditions are met, predict the location of the lane restriction on the second road;

According to the predicted position of the lane restriction on the second road, a notification or a control signal is issued to control the movement of the vehicle.

24.

A device that can be installed on a vehicle, the device comprising:

A camera adapted to capture multiple images of a road within the field of view of the camera when the camera is mounted on the vehicle and the vehicle is in motion, and;

The processor, which is configured as:

Obtaining the plurality of images from the camera;

Applying image processing to the plurality of images to identify one or more indicators of the construction area in the plurality of images, and

According to the one or more indicators identified in the construction area, a signal indicating that the vehicle is approaching the construction area is output.

25.

The device of claim 24, wherein the processor is configured to:

Based on the one or more indicators, it is determined that the vehicle is approaching the construction area, and

According to the determination that the vehicle is approaching the construction area, a signal indicating that the vehicle is approaching the construction area is output.

26.

The device of claim 24, wherein the output signal is perceivable by a driver of the vehicle.

27.

The device of claim 24, wherein the output signal is a control signal for controlling the movement of the vehicle.

28.

The device according to claim 24, wherein the one or more indicators of the construction zone include one of the following: a cone existing in the vicinity of the road, a cylinder existing in the vicinity of the road , Poles existing in the vicinity of the road, and obstacles existing in the road.

29.

The apparatus of claim 24, wherein the one or more indicators of the construction area include predetermined traffic signs.

30.

The apparatus of claim 24, wherein the one or more indicators of the construction area include a lane line of the construction area.

31.

The device of claim 24, wherein the instructions cause the system to:

According to the determination that the vehicle is approaching the construction area, the moving vehicle is moved according to the indicator of the construction area.

32.

The device according to claim 31, wherein the moving the moving vehicle according to the indicator of the construction zone comprises one of the following: reducing the speed of the vehicle; changing the lane in which the vehicle travels; realizing Avoid cones, cylinders, poles, and/or obstacles; and cause to follow the lane restrictions in the construction area instead of the traditional lane restrictions.

33.

The device of claim 24, wherein the instructions cause the system to:

Based on the one or more indicators, it is determined that the construction area is a complex construction area, and

According to the determination that the construction area is a complex construction area, the automatic vehicle control function is disabled.

34.

The apparatus of claim 33, wherein determining that the construction area is a complex construction area includes identifying one of the following: a cone existing in the lane where the vehicle is located, and presence in the lane where the vehicle is located The cylindrical body of the vehicle, the rod that exists in the lane where the vehicle is located, and the obstacles that exist in the lane where the vehicle is located.

35.

The apparatus of claim 33, wherein determining that the construction area is a complex construction area includes identifying that the number of cones, cylinders, poles, and/or obstacles in the vicinity of the vehicle exceeds a predetermined threshold number.

36.

IPR2025-00943
Tesla EX1002 Page 1090

A computer system including:

Memory for storing instructions, and

A processor, the processor executes the instructions to cause the system:

Obtaining a plurality of images from a camera, the camera being adapted to capture an image of a road in the field of view of the camera when the camera is installed on the vehicle and the vehicle is in motion;

Applying image processing to the plurality of images to identify one or more indicators of the construction area in the plurality of images, and

According to the one or more indicators identified in the construction area, a signal indicating that the vehicle is approaching the construction area is output.

37.

One method includes:

At a system including a memory storing instructions and a processor executing the instructions:

Obtaining a plurality of images from a camera, the camera being adapted to capture an image of a road in the field of view of the camera when the camera is mounted on a vehicle and the vehicle is in motion;

Applying image processing to the plurality of images to identify one or more indicators of the construction area in the plurality of images, and

According to the one or more indicators identified in the construction area, a signal indicating that the vehicle is approaching the construction area is output.

38.

A vehicle including:

main body;

A camera adapted to capture multiple images of a road within the field of view of the camera when the camera is mounted on the vehicle and the vehicle is in motion, and;

The processor, which is configured as:

Obtaining the plurality of images from the camera;

Applying image processing to the plurality of images to identify one or more indicators of the construction area in the plurality of images, and

According to the one or more indicators identified in the construction area, a signal indicating that the vehicle is approaching the construction area is output.

(19)中华人民共和国国家知识产权局

(12)发明专利申请

(10)申请公布号 CN 107111742 A
(43)申请公布日 2017.08.29

(21)申请号 201580055136.5

(22)申请日 2015.08.17

(30)优先权数据
62/038,426 2014.08.18 US

(85)PCT国际申请进入国家阶段日
2017.04.11

(86)PCT国际申请的申请数据
PCT/IL2015/050833 2015.08.17

(87)PCT国际申请的公布数据
WO2016/027270 EN 2016.02.25

(71)申请人 无比视视觉技术有限公司
地址 以色列耶路撒冷

(72)发明人 希拉·阿哈龙尼 玛丽亚·迪谢尔 阿米尔·哈雷尔 马泽尔·拉巴尼 沙克德·沙玛 科比·霍恩

(74)专利代理机构 北京安信方达知识产权代理有限公司 11262
代理人 陆建萍 郑霞

(51)Int.Cl.
G06K 9/00(2006.01)

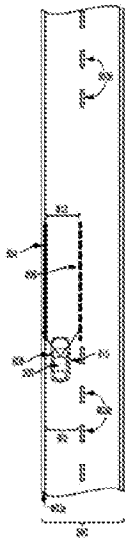
权利要求书4页 说明书36页 附图33页

(54)发明名称

在导航中对车道限制和施工区域的识别和预测

(57)摘要

公开了使用相机以提供自主和/或驾驶者辅助导航特征的系统和方法。在一些实施中，提供了用于预测第一道路车道限制的位置的技术。该系统可接收在车辆附近的道路的多个图像，识别第一道路车道限制，并且，当确定车道预测条件被满足时，预测第二道路车道限制的位置。在一些实施例中，提供了用于检测施工区并对其做出响应的技术。该系统可接收在车辆附近的道路的多个图像，识别图像中的施工区的指示物，确定车辆接近施工区，并输出指示车辆接近施工区的信号。



CN 107111742 A

IPR2025-00943
Tesla EX1002 Page 1092

1.一种能够安装在车辆上的设备，所述设备包括：

相机，所述相机在安装在所述车辆上且所述车辆处于运动状态下时，适于捕获所述相机的视场内的道路的多个图像；

处理器，其被配置为：

从所述相机获得所述多个图像；

对所述多个图像应用图像处理，以识别在所述多个图像中的至少一个图像中出现的第一道路车道限制；

基于所述多个图像，确定一个或多个界定的车道预测条件被满足；

根据一个或多个界定的车道预测条件被满足的确定，预测第二道路车道限制的位置；

根据所述第二道路车道限制的所预测的位置，发布通知或控制信号，以控制所述车辆的运动。

2.根据权利要求1所述的设备，其中，所述处理器被配置为根据所述第一道路车道限制的位置，发布通知或控制信号，以引导所述车辆的运动。

3.根据权利要求1所述的设备，其中，所述处理器被配置为发布通知或控制信号，以引导所述车辆沿着所述第一道路车道限制和所述第二道路车道限制之间的车道或区域运动。

4.根据权利要求1所述的设备，其中，所述处理器被配置为发布用于移动所述车辆的通知或控制信号，发布用于引导所述车辆的运动的通知或控制信号，以避开所述第一道路车道限制和所述第二道路车道限制。

5.如权利要求1所述的设备，其中，所述一个或多个界定的车道预测条件包括随着所述车辆在所述道路上移动，基于所述多个图像，所述第二道路车道限制停止被识别。

6.如权利要求1所述的设备，其中，所述一个或多个界定的车道预测条件包括所述第一道路车道限制被识别了至少阈值时间或阈值距离。

7.如权利要求1所述的设备，其中，所述一个或多个界定的车道预测条件包括护栏存在于所述车辆的附近。

8.如权利要求7所述的设备，其中，确定护栏存在于所述车辆的附近包括在所述一个或多个图像中检测由支撑所述护栏的柱创建的图案。

9.如权利要求7所述的设备，其中，所述一个或多个界定的车道预测条件包括所述护栏将阴影投射到所述道路上。

10.如权利要求1所述的设备，其中，所述一个或多个界定的车道预测条件包括在所述多个图像中检测到在所述道路上的车道的一侧上的多个基本上平行的边缘。

11.如权利要求1所述的设备，其中，预测所述第二道路车道限制的位置包括考虑所述第一道路车道限制的位置。

12.如权利要求1所述的设备，其中，预测所述第二道路车道限制的位置包括基于有关车道的先前确定的宽度的信息预测所述车道的宽度。

13.如权利要求1所述的设备，其中，所述处理器被配置为：

确定一个或多个界定的车道预测取消条件是否满足，以及

根据一个或多个界定的车道预测取消条件被满足的确定，停止预测所述第二道路车道限制的位置。

14.如权利要求13所述的设备，其中，所述一个或多个界定的车道预测取消条件包括，

IPR2025-00943
Tesla EX1002 Page 1093

在停止被识别之后，所述第二道路车道限制基于所述一个或多个图像被识别。

15.如权利要求13所述的设备，其中，所述一个或多个界定的车道预测取消条件包括驾驶者在车道之间移动所述车辆。

16.如权利要求13所述的设备，其中，所述一个或多个界定的车道预测取消条件包括所述车辆横穿超过预定的锐度阈值的曲线。

17.如权利要求13所述的设备，其中，确定一个或多个界定的车道预测取消条件是否被满足包括：

选择最大车道预测距离设置，以及

确定在预测所述第二道路车道限制的位置时所述车辆行进的距离是否大于与所述最大车道预测距离设置相关联的最大车道预测距离。

18.如权利要求17所述的设备，其中，选择最大车道预测距离设置包括，根据第一车道与第二车道在所述车辆的附近合并的确定，选择最大车道预测距离设置，使得所述最大车道预测距离大于默认的最大车道预测距离。

19.如权利要求17所述的设备，其中，选择最大车道预测距离设置包括，根据界定的护栏条件被满足的确定，选择没有最大车道预测距离的最大车道预测距离设置。

20.如权利要求17所述的设备，其中，选择最大车道预测距离设置包括，根据界定的护栏条件被满足的确定，选择最大车道预测距离设置，使得所述最大车道预测距离大于默认的最大车道预测距离。

21.一种计算机系统，包括：

储存指令的存储器，以及

处理器，所述处理器执行所述指令以使系统：

从相机获得多个图像，所述相机当被安装在所述车辆上且所述车辆处于运动状态下时适于捕获所述相机的视场内的道路的图像；

对所述多个图像应用图像处理，以识别在所述多个图像中的至少一个图像中出现的第一道路车道限制；

基于所述多个图像，确定一个或多个界定的车道预测条件被满足；

根据一个或多个界定的车道预测条件被满足的确定，预测第二道路车道限制的位置；

根据所述第二道路车道限制的所预测的位置，发布通知或控制信号，以引导所述车辆的运动。

22.一种方法，包括：

在包括储存指令的存储器和执行所述指令的处理器的系统处：

从相机获得多个图像，所述相机当被安装在车辆上且所述车辆处于运动状态下时适于捕获所述相机的视场内的道路的图像；

对所述多个图像应用图像处理，以识别在所述多个图像中的至少一个图像中出现的第一道路车道限制；

基于所述多个图像，确定一个或多个界定的车道预测条件被满足；

根据一个或多个界定的车道预测条件被满足的确定，预测第二道路车道限制的位置；

根据所述第二道路车道限制的所预测的位置，发布通知或控制信号，以引导所述车辆的运动。

IPR2025-00943
Tesla EX1002 Page 1094

23.一种车辆,包括:

主体;

相机,所述相机当被安装在所述车辆上且所述车辆处于运动状态下时适于捕获在所述相机的视场内的道路的多个图像;

处理器,其被配置为:

从所述相机获取多个图像;

对所述多个图像应用图像处理,以识别在所述多个图像中的至少一个图像中出现的第一道路车道限制;

基于所述多个图像,确定一个或多个界定的车道预测条件被满足;

根据一个或多个界定的车道预测条件被满足的确定,预测第二道路车道限制的位置;

根据所述第二道路车道限制的所预测的位置,发布通知或控制信号,以控制所述车辆的运动。

24.一种能够安装在车辆上的设备,所述设备包括:

相机,所述相机当被安装在所述车辆上且所述车辆处于运动状态下时适于捕获在所述相机的视场内的道路的多个图像,以及;

处理器,其被配置为:

从所述相机获得所述多个图像;

对所述多个图像应用图像处理,以在所述多个图像中识别施工区的一个或多个指示物,以及

根据识别到施工区的所述一个或多个指示物,输出指示所述车辆接近施工区的信号。

25.如权利要求24所述的设备,其中,所述处理器被配置为:

基于所述一个或多个指示物,确定所述车辆接近施工区,以及

根据所述车辆接近施工区的所述确定,输出指示所述车辆接近施工区的信号。

26.如权利要求24所述的设备,其中,所述输出信号能够由所述车辆的驾驶者察觉到。

27.如权利要求24所述的设备,其中,所述输出信号是控制所述车辆的运动的控制信号。

28.如权利要求24所述的设备,其中,施工区的所述一个或多个指示物包括以下中的一个:存在于所述道路的附近的圆锥体、存在于所述道路的附近的圆筒体、存在于所述道路的附近的杆、以及存在于所述道路中的障碍物。

29.如权利要求24所述的设备,其中,施工区的所述一个或多个指示物包括预定的交通标志。

30.如权利要求24所述的设备,其中,施工区的所述一个或多个指示物包括施工区车道线。

31.如权利要求24所述的设备,其中,所述指令使所述系统:

根据所述车辆接近施工区的所述确定,使移动的车辆根据所述施工区的指示物移动。

32.如权利要求31所述的设备,其中,使所述移动的车辆根据所述施工区的指示物移动包括以下中的一个:使所述车辆的速度降低;使所述车辆行进的车道改变;实现避开圆锥体、圆筒体、杆和/或障碍物;以及引起遵循施工区车道限制而不是传统的车道限制。

33.如权利要求24所述的设备,其中,所述指令使所述系统:

IPR2025-00943
Tesla EX1002 Page 1095

基于所述一个或多个指示物，确定所述施工区是复杂的施工区，以及

根据所述施工区是复杂的施工区的确定，禁用自动车辆控制功能。

34.如权利要求33所述的设备，其中，确定所述施工区是复杂的施工区包括识别以下中的一个：存在于所述车辆所在的车道中的圆锥体、存在于所述车辆所在的车道中的圆筒体、存在于所述车辆所在的车道中的杆、以及存在于所述车辆所在的车道中的障碍物。

35.如权利要求33所述的设备，其中，确定所述施工区是复杂的施工区包括识别在所述车辆附近的圆锥体、圆筒体、杆和/或障碍物的数量超过预定的阈值数量。

36.一种计算机系统，包括：

储存指令的存储器，以及

处理器，所述处理器执行所述指令以使所述系统：

从相机获得多个图像，所述相机在被安装在所述车辆上且所述车辆处于运动状态下时适于捕获在所述相机的视场内的道路的图像；

对所述多个图像应用图像处理，以在所述多个图像中识别施工区的一个或多个指示物，以及

根据识别到施工区的所述一个或多个指示物，输出指示所述车辆接近施工区的信号。

37.一种方法，包括：

在包括储存指令的存储器和执行所述指令的处理器的系统处：

从相机获得多个图像，所述相机在被安装在车辆上且所述车辆处于运动状态下时适于捕获在所述相机的视场内的道路的图像；

对所述多个图像应用图像处理，以在所述多个图像中识别施工区的一个或多个指示物，以及

根据识别到施工区的所述一个或多个指示物，输出指示所述车辆接近施工区的信号。

38.一种车辆，包括：

主体；

相机，所述相机在被安装在所述车辆上且所述车辆处于运动状态下时适于捕获在所述相机的视场内的道路的多个图像，以及；

处理器，其被配置为：

从所述相机获得所述多个图像；

对所述多个图像应用图像处理，以在所述多个图像中识别施工区的一个或多个指示物，以及

根据识别到施工区的所述一个或多个指示物，输出指示所述车辆接近施工区的信号。

IPR2025-00943
Tesla EX1002 Page 1096

在导航中对车道限制和施工区域的识别和预测

[0001] 相关申请的交叉引用

[0002] 本申请要求美国临时专利申请号62/038,426的权益，其公开内容通过引用并入本文。

[0003] 公开领域

[0004] 本公开大体上涉及自主驾驶和/或驾驶者辅助技术，特别是涉及使用相机以提供自主驾驶和/或驾驶者辅助技术的特征的系统和方法。

[0005] 背景

[0006] 随着技术不断地进步，能够在道路上进行导航的全自主车辆的目标很快就能实现。首先，自主车辆可在没有人类操作者的输入的情况下能够辨认其环境并导航。自主车辆也可考虑各种因素并基于这些因素做出适当的决策，以安全地并准确地到达预期的目的地。例如，在车辆通常行进在道路上时会遇到各种对象，诸如其他车辆和行人。自主驾驶系统可识别车辆的环境中的这些对象，并采取适当和及时的行动来避免碰撞。另外，自主驾驶系统可辨认调节车辆移动(例如，车辆何时必须停止和可以走、车辆必须不能超过的速度、车辆必须位于道路上的哪里等)的其他指示物，诸如交通信号、交通标志标志和车道标记。自主驾驶系统可能需要确定车辆何时应改变车道、在交叉口转弯、改变道路等。从这些示例明显看出的是，为了提供能够安全和准确导航的自主车辆，许多因素可能需要加以解决。

[0007] 概述

[0008] 当前的自主或驾驶者辅助(例如，部分自主)的车辆系统在车道线(例如，划分车道的标记线)或其他标记没有清楚指示道路上的车道限制的位置时，可能无法充分辨认车道。例如，当前的系统在车道限制不能从在接收的道路图像中所检测的线或标记中被辨认出来时可能不能正常工作；或者，当前的系统在车道限制在现实中不存在的情况下可能不能准确辨认车道限制。当前的自主或驾驶者辅助的车辆系统也可能无法充分辨认或响应于道路上的施工区，在施工区中可能要求修改的车辆控制，以及在施工区中以不同方式响应于某些道路元素可能是有利的。发明人已经认识到这些和其他问题，这些和其他问题通过本公开将变得明显，并通过本文中所公开的技术得以解决。

[0009] 与本公开一致的一些实施例提供了用于自主或驾驶者辅助的车辆导航的系统和方法。一些公开的实施例可使用相机来提供自主车辆导航特征。例如，与一些公开的实施例一致的被公开的系统可包括监控车辆的环境的一个、两个或多个相机并基于由一个或多个相机捕获的图像的分析引起导航响应。

[0010] 根据一些实施例，计算机系统包括储存指令的存储器和执行指令以使系统能够执行以下操作的处理器：从相机接收移动的车辆附近的道路的多个图像；基于多个图像识别第一道路车道限制；基于多个图像确定一个或多个界定的车道预测条件被满足；根据一个或多个界定的车道预测条件被满足的确定，预测第二道路车道限制的位置；以及使移动的车辆能够避免第一道路车道限制和第二道路车道限制。

[0011] 根据一些实施例，计算机系统包括储存指令的存储器和执行指令以使系统能够执行以下操作的处理器：从相机接收移动的车辆附近的道路的多个图像；在多个图像中识别

IPR2025-00943
Tesla EX1002 Page 1097

施工区的一个或多个指示物;基于一个或多个指示物确定车辆接近施工区;以及,根据车辆接近施工区的确定,输出指示车辆接近施工区的信号。

[0012] 附图简述

[0013] 图1是与公开的实施例一致的示例性系统的图形表示。

[0014] 图2A是包括与公开的实施例一致的系统的示例性车辆的图形侧视图表示。

[0015] 图2B是与公开的实施例一致的图2A中示出的车辆和系统的图形俯视图表示。

[0016] 图2C是包括与公开的实施例一致的系统的另一车辆实施例的图形俯视图表示。

[0017] 图2D是包括与公开的实施例一致的系统的又一车辆实施例的图形俯视图表示。

[0018] 图2E是包括与公开的实施例一致的系统的又一车辆实施例的图形俯视图表示。

[0019] 图2F是与公开的实施例一致的示例性车辆控制系统的图形表示。

[0020] 图3A是与公开的实施例一致的包括后视镜和用于车辆成像系统的用户接口的车辆内部的图形表示。

[0021] 图3B是与公开的实施例一致的被配置为定位在后视镜之后并对着车辆挡风玻璃的相机底座的示例的图示。

[0022] 图3C是与公开的实施例一致的图3B中示出的相机底座从不同角度的图示。

[0023] 图3D是与公开的实施例一致的被配置为定位在后视镜之后并对着车辆挡风玻璃的相机底座的示例的图示。

[0024] 图4是与公开的实施例一致的被配置为储存用于执行一个或多个操作的指令的存储器的示例性框图。

[0025] 图5A是示出了与公开的实施例一致的用于基于单目图像分析引起一个或多个导航响应的示例性过程的流程图。

[0026] 图5B是示出了与公开的实施例一致的用于在一组图像中检测一个或多个车辆和/或行人的示例性过程的流程图。

[0027] 图5C是示出了与公开的实施例一致的用于在一组图像中检测路标记和/或车道几何形状信息的示例性过程的流程图。

[0028] 图5D是示出了与公开的实施例一致的用于在一组图像中检测交通灯的示例性过程的流程图。

[0029] 图5E是示出了与公开的实施例一致的用于基于车辆路径引起一个或多个导航响应的示例性过程的流程图。

[0030] 图5F是示出了与公开的实施例一致的用于确定领先的车辆是否正改变车道的示例性过程的流程图。

[0031] 图6是示出了与公开的实施例一致的用于基于立体图像分析引起一个或多个导航响应的示例性过程的流程图。

[0032] 图7是示出了与公开的实施例一致的用于基于三组图像的分析引起一个或多个导航响应的示例性过程的流程图。

[0033] 图8A-图8C是根据一些实施例的沿着道路行进的车辆的描绘。

[0034] 图9A-图9C是根据一些实施例的沿着道路行进的车辆的描绘。

[0035] 图10A是来自与根据一些实施例的系统相关联的相机的透视图的描绘。

[0036] 图10B是根据一些实施例的沿着道路行进的车辆的描绘。

IPR2025-00943
Tesla EX1002 Page 1098

[0037] 图11是与一些公开的实施例一致的示例性系统的存储器的图形表示。

[0038] 图12是根据一些实施例的示例性过程的流程图。

[0039] 图13是根据一些实施例的沿着道路行进的车辆的描绘。

[0040] 图14是与一些公开的实施例一致的示例性系统的存储器的图形表示。

[0041] 图15是根据一些实施例的示例性过程的流程图。

[0042] 详细描述

[0043] 以下的详细描述参照附图。只要有可能，相同的参考数字用在附图和以下的描述中，以指代相同或类似的部分。虽然本文中对几个说明性实施例进行了描述，但修改、适应和其他实施也是可能的。例如，可对附图中图示的组分进行替换、添加或修改，并且本文中所描述的说明性方法可通过将步骤替换、重排、移除或添加到公开的方法来修改。因此，以下的详细描述不限于公开的实施例和示例。

[0044] 公开的实施例提供了使用相机来提供自主导航、车辆控制和/或驾驶者辅助技术的特征的系统和方法。驾驶者辅助技术，而不是全自主驾驶，指用于在他们车辆的导航和控制中辅助驾驶者的任何合适的技术，诸如LKA(车道保持辅助)、LDW(车道偏离报警)、车辆的加速、车辆的减速、车辆的转向、控制车辆制动、使车辆对不可避免的碰撞做出准备、和/或控制车辆操作的任何方面或在控制任何这样的方面辅助驾驶者的任何其他合适的方式。在各个实施例中，系统可包括监控车辆的环境的一个、两个或多个相机。在一些实施例中，系统可提供用于预测道路车道限制的位置的技术和/或用于检测和响应于施工区的技术。

[0045] 图1是与示例性公开的实施例一致的系统100的框图表示。系统100可包括取决于特定实施的要求的各种组件。在一些实施例中，系统100可包括处理单元110、图像采集单元120、定位传感器130、一个或多个存储器单元140、150、地图数据库160和用户接口170。处理单元110可包括一个或多个处理设备。在一些实施例中，处理单元110可包括应用处理器180、图像处理器190或任何其他合适的处理设备。类似地，图像采集单元120可包括取决于特定应用的要求的任意数量的图像采集设备和组件。在一些实施例中，图像采集单元120可包括一个或多个图像捕获设备(例如，相机)，诸如图像捕获设备122、图像捕获设备124和图像捕获设备126。系统100也可包括在通信上将处理单元110连接到图像采集设备120的数据接口128。例如，数据接口128可包括用于将由图像采集设备120采集的图像数据传输到处理单元110的任何有线和/或无线的链接。

[0046] 应用处理器180和图像处理器190二者可包括各种类型的处理设备。例如，应用处理器180和图像处理器190中任意一个或二者可包括一个或多个微处理器、预处理器(诸如图像预处理器)、图形处理器、中央处理单元(CPU)、支持电路、数字信号处理器、集成电路、存储器或适于运行应用以及适于图像处理和分析的任何其他类型的设备。在一些实施例中，应用处理器180和/或图像处理器190可包括任何类型的单核或多核处理器、移动设备微控制器、中央处理单元等。可使用各种处理设备，例如包括从诸如Intel®、AMD®等的制造商可购得的处理器，并且可包括各种架构(例如，x86处理器、ARM®等)。

[0047] 在一些实施例中，应用处理器180和/或图像处理器190可包括从Mobileye®可购得的任何EyeQ系列的处理器芯片。这些处理器设计各自包括具有本地存储器和指令集的多个处理单元。这样的处理器可包括用于从多个图像传感器接收图像数据的视频输入端，并且也可包括视频输出能力。在一个实施例中，EyeQ2®使用运行在332Mhz的90纳米微米技

IPR2025-00943
Tesla EX1002 Page 1099

术。EyeQ2®架构具有两个浮点、超线程32位RISC CPU(MIPS32® 34K®核)、五个视觉计算引擎(VCE)、三个向量微码处理器(VMP®)、Denali 64位移动DDR控制器、128位内部超音速互连、双16位视频输入端和18位视频输出端控制器、16通道DMA和几个外围设备。MIPS34K CPU管理五个VCE、三个VMP.TM和DMA、第二MIPS34K CPU和多通道DMA以及其他外围设备。五个VCE、三个VMP®和MIPS34K CPU可执行多功能包应用所需的密集视觉计算。在另一示例中，EyeQ3®可用在公开的实施例中，该EyeQ3®是第三代处理器并且比EyeQ2®强大六倍。

[0048] 虽然图1描绘了被包括在处理单元110中的两个单独的处理设备，但可使用更多或更少的处理设备。例如，在一些实施例中，单一处理设备可用于完成应用处理器180和图像处理器190的任务。在其他实施例中，这些任务可由两个以上的处理设备执行。

[0049] 处理单元110可包括各种类型的设备。例如，处理单元110可包括各种设备，诸如控制器、图像预处理器、中央处理单元(CPU)、支持电路、数字信号处理器、集成电路、存储器或用于图像处理和分析的任何其他类型的设备。图像预处理器可包括用于捕获、数字化和处理来自图像传感器的图像的视频处理器。CPU可包括任意数量的微控制器或微处理器。支持电路可以是通常在本领域中众所周知的任意数量的电路，包括高速缓冲存储器、电源、时钟和输入-输出电路。存储器可储存软件，该软件在由处理器执行时控制系统的操作。存储器可包括数据库和图像处理软件。存储器可包括任意数量的随机存取存储器、只读存储器、闪存、盘驱动、光储存器、带储存器、可移除储存器和其他类型的储存器。在一个实例中，存储器可与处理单元110分开。在另一实例中，存储器可集成到处理单元110中。

[0050] 每个存储器140、150可包括软件指令，该软件指令在由处理器(例如，应用处理器180和/或图像处理器190)执行时可控制系统100的各个方面的操作。这些存储器单元可包括各种数据库和图像处理软件。存储器单元可包括随机存取存储器、只读存储器、闪存、盘驱动、光储存器、带储存器、可移除储存器和/或任何其他类型的储存器。在一些实施例中，存储器单元140、150可与应用处理器180和/或图像处理器190分开。在其他实施例中，这些存储器单元可集成到应用处理器180和/或图像处理器190中。

[0051] 定位传感器130可包括适于确定与系统100的至少一个组件相关联的位置的任意类型的设备。在一些实施例中，定位传感器130可包括GPS接收器。这样的接收器可通过处理由全球定位系统卫星传播的信号确定用户的定位和速度。来自定位传感器130的定位信息可供应用处理器180和/或图像处理器190使用。

[0052] 用户接口170可包括适于向系统100的一个或多个用户提供信息或适于接收来自系统100的一个或多个用户的输入的任何设备。在一些实施例中，用户接口170可包括用户输入设备，例如包括触摸屏、麦克风、键盘、指针设备、轨道轮(track wheels)、相机、旋钮、按钮等。借助这样的输入设备，用户可通过键入指令或信息、提供语音命令、使用按钮、指针或眼睛跟踪能力在屏幕上选择菜单选项，或者通过用于将信息传递到系统100的任何其他合适的技术，向系统100提供信息输入或命令。

[0053] 用户接口170可配备有一个或多个处理设备，该一个或多个处理设备被配置为向用户提供信息或接收来自用户的信息，并对该信息进行处理以例如供应用处理器180使用。在一些实施例中，这样的处理设备可执行用于识别和跟踪眼睛移动、接收和解释语音命令、

IPR2025-00943
Tesla EX1002 Page 1100

识别和解释触摸屏上做出的触摸和/或手势、响应于键盘输入或菜单选择等的指令。在一些实施例中，用户接口170可包括显示器、扬声器、触觉设备和/或用于向用户提供输出信息的任何其他设备。

[0054] 地图数据库160可包括用于储存对系统100有用的地图数据的任意类型的数据库。在一些实施例中，地图数据库160可包括与各种项目在参考坐标系中的定位有关的数据，该各种项目包括路、水文特征、地理特征、商行、兴趣点、餐馆、加油站等。地图数据库160不仅可储存这样的项目的位置，还可储存与这些项目有关的描述符，例如包括与任何储存的特征相关联的名称。在一些实施例中，地图数据库160可用系统100的其他组件来物理定位。可替代地或另外地，地图数据库160或其部分可相对于系统100的其他组件(例如，处理单元110)来远程定位。在这样的实施例中，来自地图数据库160的信息可通过与网络的有线或无线的数据连接(例如，通过蜂窝网络和/或互联网等)来下载。

[0055] 图像捕获设备122、124和126可各自包括适于从环境捕获至少一个图像的任意类型的设备。此外，任意数量的图像捕获设备可用于采集输入到图像处理器的图像。一些实施例可包括仅单一图像捕获设备，而其他实施例可包括两个、三个或者甚至四个或更多的图像捕获设备。以下将参照图2B-2E进一步对图像捕获设备122、124和126进行描述。

[0056] 系统100或其各种组件可被纳入到各种不同的平台。在一些实施例中，如图2A中所示，系统100可被包括在车辆200上。例如，如以上根据图1所描述的，车辆200几可配备有处理单元110和系统100的任意其他组件。虽然在一些实施例中，车辆200可仅配备有单一图像捕获设备(例如，相机)，但在诸如结合图2B-图2E那些所讨论的其他实施例中，可使用多个图像捕获设备。例如，如图2A中所示，车辆200的图像捕获设备122和124中任意一个可以是ADAS(高级驾驶者辅助系统)成像装置的部分。

[0057] 作为图像采集单元120的部分的被包括在车辆200上的图像捕获设备可位于任何合适的位置。在一些实施例中，如图2A-图2E和图3A-图3C中所示，图像捕获设备122可位于后视镜附近。该位置可提供类似于车辆200驾驶者的视线的视线，这可有助于确定什么对于驾驶者是可见和不可见的。图像捕获设备122可被定位在后视镜附近的任何位置，但将图像捕获设备122放置在镜子的驾驶者侧上可进一步有助于获得表示驾驶者的视场和/或视线的图像。

[0058] 也可使用图像采集单元120的图像捕获设备的其他位置。例如，图像捕获设备124可位于车辆200的保险杠(bumper)上或其中。这样的位置可能尤其适于具有宽视场的图像捕获设备。位于保险杠的图像捕获设备的视线可不同于驾驶者的视线，因此，保险杠图像捕获设备和驾驶者可能不会总是见到相同的对象。图像捕获设备(例如，图像捕获设备122、124和126)也可位于其他位置中。例如，图像捕获设备可位于车辆200的侧视镜中的一个或二者的上面或其中、车辆200的车顶上、车辆200的引擎罩上、车辆200的后备箱上、车辆200的侧面上、安装在车辆200的任意窗户上、定位在车辆200的任意窗户的后面或前面、以及安装在车辆200的前面和/或后面的照明轮廓中或附近。

[0059] 除了图像捕获设备之外，车辆200还可包括系统100的各种其他组件。例如，处理单元110可与车辆的引擎控制单元(ECU)集成或与其分开地被包括在车辆200上。车辆200也可配备有诸如GPS接收器的定位传感器130，并且也可包括地图数据库160以及存储器单元140和150。

IPR2025-00943
Tesla EX1002 Page 1101

[0060] 图2A是与公开的实施例一致的示例性车辆成像系统的图形侧视图表示。图2B是图2A中所示的实施例的图形俯视图图示。如图2B中所图示的，公开的实施例可包括车辆200，该车辆200在其本体中包括具有第一图像捕获设备122、第二图像捕获设备124和处理单元110的系统100，该第一图像捕获设备122被定位在后视镜附近和/或接近车辆200的驾驶者，该第二图像捕获设备124被定位在车辆200的保险杠区域(例如，保险杠区域210中的一个)上或其中。

[0061] 如图2C中所图示的，图像捕获设备122和124都可被定位在车辆200的后视镜附近和/或接近车辆200的驾驶者。另外，虽然图2B和图2C中示出了两个图像捕获设备122和124，但应理解的是，其他实施例可包括两个以上的图像捕获设备。例如，在图2D和图2E中示出的实施例中，第一图像捕获设备、第二图像捕获设备和第三图像捕获设备122、124和126被包括在车辆200的系统100中。

[0062] 如图2D中所图示的，图像捕获设备122可被定位在车辆200的后视镜附近和/或接近车辆200的驾驶者，并且图像捕获设备124和126可被定位在车辆200的保险杠区(例如，保险杠区210中的一个)上或其中。如图2E中所示，图像捕获设备122、124和126可被定位在车辆200的后视镜附近和/或接近车辆200的驾驶者。所公开的实施例不限于任何特定数量的图像捕获设备和其配置，并且图像捕获设备可被定位在车辆200之内和/或之上的任何适当的位置。

[0063] 应理解的是，所公开的实施例不限于车辆并且能够在其他背景下应用。还应理解的是，所公开的实施例不限于特定类型的车辆200，并且可适用于所有类型的车辆，包括汽车、卡车、拖车、摩托车、自行车、自平衡运输设备和其他类型的车辆。

[0064] 第一图像捕获设备122可包括任何合适类型的图像捕获设备。图像捕获设备122可包括光轴。在一个实例中，图像捕获设备122可包括具有全局快门的Aptina M9V024WVGA传感器。在其他实施例中，图像捕获设备122可提供1280x960像素的分辨率，并可包括卷帘快门。图像捕获设备122可包括各种光学元件。在一些实施例中，可包括一个或多个透镜，例如，以提供对于图像捕获设备期望的焦距和视场。在一些实施例中，图像捕获设备122可与6mm透镜或12mm透镜相关联。在一些实施例中，图像捕获设备122可被配置为捕获具有如图2D中所图示的期望的视场(FOV)202的图像。例如，图像捕获设备122可被配置为具有规则的FOV，诸如在40度到56度的范围内，包括46度FOV、50度FOV、52度FOV或更大的FOV。可替代地，图像捕获设备122可被配置为具有23度到40度的范围内的窄FOV，诸如28度FOV或36度FOV。另外，图像捕获设备122可被配置为具有100度到180度的范围内的广FOV。在一些实施例中，图像捕获设备122可包括广角保险杠相机或一个高达180度FOV的广角保险杠相机。

[0065] 第一图像捕获设备122可采集相对于与车辆200相关联的场景的多个第一图像。多个第一图像中的每个可作为一系列图像扫描线被采集，该多个第一图像可使用卷帘快门来捕获。每个扫描线可包括多个像素。

[0066] 第一图像捕获设备122可具有与第一系列图像扫描线中的每个的采集相关联的扫描速率。扫描速率可指图像传感器可采集与包括在特定扫描线中的每个像素相关联的图像数据的速率。

[0067] 图像捕获设备122、124和126可包含任何合适类型和数量的图像传感器，例如包括CCD传感器或CMOS传感器。在一个实施例中，CMOS图像传感器可连同卷帘快门来使用，使得

IPR2025-00943
Tesla EX1002 Page 1102

成行的每个像素一次读取一个，并且行的扫描在逐行的基础上进行，直至整个图像帧被捕获。在一些实施例中，行可相对于帧从上到下被依次捕获。

[0068] 卷帘快门的使用可能导致不同行的像素在不同时间被曝光和捕获，这可能在所捕获的图像帧中引起歪斜(skew)和其他图像伪影(artifact)。另一方面，当图像捕获设备122被配置为与全局快门或同步快门一起操作时，所有像素可在相同的时间量内并在共同的曝光时段期间被曝光。因此，从采用全局快门的系统所收集的帧中的图像数据表示在特定时间的整个FOV(诸如FOV 202)的快照。相反，在卷帘快门应用中，帧中的每行在不同时间被曝光，并且数据在不同时间被捕获。因此，移动的对象在具有卷帘快门的图像捕获设备中可能出现扭曲(distort)。以下将对该现象更详细地进行描述。

[0069] 第二图像捕获设备124和第三图像捕获设备126可以是任意类型的图像捕获设备。类似于第一图像捕获设备122，图像捕获设备124和126中的每个可包括光轴。在一个实例中，图像捕获设备124和126中的每个可包括具有全局快门的Aptina M9V024WVGA传感器。可替代地，图像捕获设备124和126中的每个可包括卷帘快门。类似于图像捕获设备122，图像捕获设备124和126可被配置为包括各种透镜和光学元件。在一些实施例中，与图像捕获设备124和126相关联的透镜可提供FOV(诸如FOV 204和206)，该FOV同与图像捕获设备122相关联的FOV(诸如FOV 202)相同或不同(例如，更窄)。例如，图像捕获设备124和126可具有40度、30度、26度、23度、20度或更低的FOV。

[0070] 图像捕获设备124和126可采集相对于与车辆200相关联的场景的多个第二图像和第三图像。多个第二图像和第三图像中的每个可作为第二系列图像扫描线和第三系列的图像扫描线被采集，该多个第二图像和第三图像中每个可使用卷帘快门来捕获。每个扫描线或扫描行可具有多个像素。图像捕获设备124和126可具有与包括在第二系列和第三系列中的每个图像扫描线的采集相关联的第二扫描速率和第三扫描速率。

[0071] 每个图像捕获设备122、124和126可被定位在相对于车辆200的任何合适的位置和定向处。可选择图像捕获设备122、124和126的相对定向，以有助于将从图像捕获设备采集的信息融合在一起。例如，在一些实施例中，与图像捕获设备124相关联的FOV(诸如FOV 204)可与图像捕获设备122相关联的FOV(诸如FOV 202)和与图像捕获设备126相关联的FOV(诸如FOV 206)部分或全部重叠。

[0072] 图像捕获设备122、124和126可位于车辆200上的任何合适的相对高度上。在一个实例中，在图像捕获设备122、124和126之间有高度差，这可提供足够的视差信息以实现立体分析。例如，如图2A中所示，两个图像捕获设备122和124处于不同的高度。在图像捕获设备122、124和126之间也可存在横向位移差，例如，通过处理单元110给出用于立体分析的额外的视差信息。如图2C和图2D中所示，横向位移上的差可表示为dx。在一些实施例中，前位移和后位移(例如，范围位移)可存在于图像捕获设备122、124和126之间。例如，图像捕获设备122可位于图像捕获设备124和/或图像捕获设备126之后的0.5到2米或更多处。这种类型的位移可使图像捕获设备中的一个能够覆盖其他图像捕获设备的潜在盲点。

[0073] 图像捕获设备122可具有任何合适的分辨率能力(例如，与图像传感器相关联的像素的数量)，并且与图像捕获设备122相关联的图像传感器的分辨率可高于、低于与图像捕获设备124和126相关联的图像传感器的分辨率，或与其相同。在一些实施例中，与图像捕获设备122和/或图像捕获设备124和126相关联的图像传感器可具有640x480、1024x768、

IPR2025-00943
Tesla EX1002 Page 1103

1280x960的分辨率或任何其他合适的分辨率。

[0074] 帧速率(例如,图像捕获设备在移动以捕获与下个图像帧相关联的像素数据之前,采集一个图像帧的一组像素数据的速率)可以是可控的。与图像捕获设备122相关联的帧速率可高于、低于与图像捕获设备124和126相关联的帧速率,或与其相同。与图像捕获设备122、124和126相关联的帧速率可取决于可能影响帧速率的定时的各种因素。例如,一个或多个图像捕获设备122、124和126可包括可选择的像素延迟时段,该可选择的像素延迟时段在采集与图像捕获设备122、124和/或126中图像传感器的一个或多个像素相关联的图像数据之前或之后被应用。通常,对应于每个像素的图像数据可根据设备的时钟速率(例如,每个时钟周期一个像素)来进行采集。另外,在包括卷帘快门的实施例中,一个或多个图像捕获设备122、124和126可包括可选择的水平消隐时段,该可选择的水平消隐时段在采集与图像捕获设备122、124和/或126中的图像传感器的像素行相关联的图像数据之前或之后被应用。此外,一个或多个图像捕获设备122、124和/或126可包括可选择的竖直消隐时段,该可选择的竖直消隐时段在采集与图像捕获设备122、124和126的图像帧相关联的图像数据之前或之后被应用。

[0075] 这些定时控制可使与图像捕获设备122、124和126相关联的帧速率即使在每个的线扫描速率不同的情况下也能够同步。另外,如以下将更详细讨论的,除了其他因素(例如,图像传感器分辨率、最大线扫描速率等)之外,这些可选择的定时控制还可使来自图像捕获设备122的FOV与图像捕获设备124和126的一个或多个FOV重叠的区域中的图像捕获,即使在图像捕获设备122的视场与图像捕获设备124和126的FOV不同的情况下,也能够同步。

[0076] 图像捕获设备122、124和126中的帧速率定时可取决于相关联的图像传感器的分辨率。例如,假设两个设备有类似的线扫描速率,如果一个设备包括具有640x480的分辨率的图像传感器,以及另一设备包括具有1280x960的分辨率的图像传感器,则从具有较高分辨率的传感器采集图像数据的帧将需要更多的时间。

[0077] 可影响图像捕获设备122、124和126中的图像数据采集的定时的另一因素是最大线扫描速率。例如,来自包括在图像捕获设备122、124和126中的图像传感器的图像数据的行的采集将需要一些最小时间量。假设没有增加像素延迟时段,用于采集图像数据的行的该最小时间量将与特定设备的最大线扫描速率有关。给予较高的最大线扫描速率的设备有可能提供比具有较低的最大线扫描速率的设备更高的帧速率。在一些实施例中,一个或多个图像捕获设备124和126可具有比与图像捕获设备122相关联的最大线扫描速率更高的最大线扫描速率。在一些实施例中,图像捕获设备124和/或126的最大线扫描速率可以是图像捕获设备122的最大线扫描速率的1.25倍、1.5倍、1.75倍或2倍或者更多倍。

[0078] 在另一实施例中,图像捕获设备122、124和126可具有相同的最大线扫描速率,但图像捕获设备122可在小于或等于其最大扫描速率的扫描速率来操作。可对系统进行配置,使得一个或多个图像捕获设备124和126在等于图像捕获设备122的线扫描速率的线扫描速率下运行。在其他实例中,可对系统进行配置,使得图像捕获设备124和/或图像捕获设备126的线扫描速率可以是图像捕获设备122的线扫描速率的1.25倍、1.5倍、1.75倍或2倍或者更多倍。

[0079] 在一些实施例中,图像捕获设备122、124和126可以是不对称的。也就是说,它们可包括具有不同的视场(FOV)和焦距的相机。例如,图像捕获设备122、124和126的视场可包括

IPR2025-00943
Tesla EX1002 Page 1104

相对于车辆200的环境的任何期望的区域。在一些实施例中，一个或多个图像捕获设备122、124和126可被配置为从车辆200的前方、车辆200后方、车辆200的侧面或它们的组合的环境中采集图像数据。

[0080] 此外，与每个图像捕获设备122、124和/或126相关联的焦距可以是可选择的(例如，通过包含适当的透镜等)，使得每个设备采集在相对于车辆200的期望距离范围处的对象的图像。例如，在一些实施例中，图像捕获设备122、124和126可采集在距车辆几米内的近距离对象的图像。图像捕获设备122、124和126也可被配置为采集在距车辆较远的范围(例如，25m、50m、100m、150m或更远)处的对象的图像。此外，可选择图像捕获设备122、124和126的焦距，使得一个图像捕获设备(例如，图像捕获设备122)可采集相对接近车辆(例如，在10m内或在20m内)的对象的图像，而其他图像捕获设备(例如，图像捕获设备124和126)可采集距车辆200更远对象(例如，大于20m、50m、100m、150m等)的图像。

[0081] 根据一些实施例，一个或多个图像捕获设备122、124和126的FOV可具有广角。例如，具有140度的FOV，特别是对于可用于捕获车辆200附近区域的图像的图像捕获设备122、124和126来说，可能是有利的。例如，图像捕获设备122可用于捕获车辆200的右侧或左侧的区域的图像，并且在这样的实施例中，可能需要图像捕获设备122具有广FOV(例如，至少140度)。

[0082] 在一些实施例中，系统可根据预定限制或动态确定的限制来界定车辆的“附近”。车辆的附近例如可认为是车辆在任何一个或多个方向上的某距离(诸如1米、5米、10米、25米、50米、100米、250米、500米或1000米)内的任何区域。在一些实施例中，被界定为车辆的附近的区域可以是不规则形状的，使得其在不同方向上延伸不同距离；由车辆的行驶所界定的前向方向上的附近例如可延伸几百米，而在后向或侧面的方向上的附近例如仅可延伸几十米或更少。在一些实施例中，界定的车辆的附近可根据由一个或多个图像捕获设备捕获的信息动态地变化，使得附近可例如根据道路的哪些部分在图像捕获设备的FOV中是可见的而变化。在一些实施例中，车辆的附近可例如根据道路或车辆的任何可察觉状况或状态动态地变化，该任何可察觉条件或状态包括车辆速率、张贴的速率限制、道路状况、道路拥堵、道路大小、能见度距离、环境亮度、时间、其他车辆的存在、行人的存在或任何其他合适的因素。

[0083] 与图像捕获设备122、124和126中的每个相关联的视场可取决于各自的焦距。例如，随着焦距增加，相应的视场减少。

[0084] 图像捕获设备122、124和126可被配置为具有任何合适的视场。在一个具体事例中，图像捕获设备122可具有46度的水平FOV，图像捕获设备124可具有23度的水平FOV，以及图像捕获设备126可具有介于23度和46度之间的水平FOV。在另一实例中，图像捕获设备122可具有52度的水平FOV，图像捕获设备124可具有26度的水平FOV，以及图像捕获设备126可具有介于26度和52度之间的水平FOV。在一些实施例中，图像捕获设备122的FOV与图像捕获设备124和/或图像捕获设备126的FOV的比可从1.5变化到2.0。在其他实施例中，该比可在1.25和2.25之间变化。

[0085] 可对系统100进行配置，使得图像捕获设备122的视场与图像捕获设备124和/或图像捕获设备126的视场至少部分或全部重叠。在一些实施例中，可对系统100进行配置，使得图像捕获设备124和126的视场例如在图像捕获设备122的视场的范围内(例如，更窄)并与

IPR2025-00943
Tesla EX1002 Page 1105

其共享共同中心。在其他实施例中，图像捕获设备122、124和126可捕获临近的FOV，或者在它们的FOV中可具有部分重叠。在一些实施例中，可使图像捕获设备122、124和126的视场对齐，使得较窄的FOV图像捕获设备124和/或126的中心可位于较广的FOV设备122的视场的下半部分。

[0086] 图2F是与公开的实施例一致的示例性车辆控制系统的图形表示。如图2F中所示，车辆200可包括节流系统220、制动系统230和转向系统240。系统100可通过一个或多个数据链路(例如，任何有线和/或无线的链路或者用于传输数据的链路)向节流系统220、制动系统230和转向系统240中的一个或多个提供输入(例如，控制信号)。例如，基于由图像捕获设备122、124和/或126采集的图像的分析，系统100可向节流系统220、制动系统230和转向系统240中的一个或多个提供控制信号，以导航车辆200(例如，通过引起加速、转弯、车道偏移等)。此外，系统100可接收来自指示车辆200的运行状况(例如，速率，车辆200是否制动和/或转弯等)的流系统220、制动系统230和转向系统240中的一个或多个的输入。以下结合图4-图7提供了进一步的细节。

[0087] 如图3A中所示，车辆200也可包括用于与驾驶者或车辆200的乘客交互的用户接口170。例如，车辆应用中的用户接口170可包括触摸屏320、旋钮330、按钮340和麦克风350。车辆200的驾驶者或乘客也可使用手柄(例如，位于车辆200的转向柱上的或其附近的，例如包括转弯信号手柄)、按钮(例如，位于车辆200的方向盘上)等，以与系统100交互。在一些实施例中，麦克风350可被定位为临近后视镜310。类似地，在一些实施例中，图像捕获设备122可位于后视镜310附近。在一些实施例中，用户接口170也可包括一个或多个扬声器360(例如，车辆音频系统的扬声器)。例如，系统100可经由扬声器360提供各种通知(例如，警报)。

[0088] 图3B-图3D是与公开的实施例一致的、被配置为定位在后视镜(例如，后视镜310)之后并对着车辆挡风玻璃的示例性相机底座370的图示。如图3B中所示，相机底座370可包括图像捕获设备122、124和126。图像捕获设备124和126可被定位在防眩板380之后，该防眩板380可与车辆挡风玻璃齐平，并包括薄膜和/或抗反射的材料的组合物。例如，防眩板380可被定位成使得其对着具有匹配斜率的车辆挡风玻璃对准。在一些实施例中，例如如图3D中所描绘的，图像捕获设备122、124和126中的每个可被定位在防眩板380之后。公开的实施例不限于图像捕获设备122、124和126、相机底座370以及防眩板380的任何特定配置。图3C是图3B中所示的相机底座370从前视角的图示。

[0089] 具有本公开的好处的本领域中的技术人员将认识到的是，可对上述公开的实施例做出多种变化和/或修改。例如，不是所有组件对于系统100的运行是必需的。此外，任何组件可位于系统100的任何适当的部分中，并且组件可被重新布置到各种配置中，同时提供公开的实施例的功能。因此，上述配置是示例，并且无论以上所讨论的配置，系统100都可提供广泛的功能以分析车辆200的周围环境，并且响应于该分析，导航和/或以其他方式控制和/或操作车辆200。车辆200的导航、控制和/或操作可包括(直接或经由诸如以上所提到的控制器的中间控制器)启用和/或禁用与车辆200相关联的各种特征、组件、设备、模式、系统和/或子系统。可替代地或另外地，导航、控制和/或操作可包括例如通过提供视觉、音频、触觉和/或其他感官警报以及/或者指示与可位于车辆200的内部或外部的用户、驾驶者、乘客、路人和/或其他车辆或用户的交互。

[0090] 如以下进一步详细讨论的且与各种公开的实施例一致的，系统100可提供与自主

IPR2025-00943
Tesla EX1002 Page 1106

驾驶和/或驾驶者辅助技术有关的各种特征。例如，系统100可分析图像数据、定位数据（例如，GPS位置信息）、地图数据、速率数据和/或来自包括在车辆200中的传感器的数据。系统100可例如从图像采集单元120、定位传感器130和其他传感器收集用于分析的数据。此外，系统100可分析所收集的数据，以确定车辆200是否应采取某种动作，并随后在没有人类干预的情况下自动采取确定的动作。例如，当车辆200在没有人类干预的情况下导航时，（例如，通过将控制信号发送给节流系统220、制动系统230和转向系统240中的一个或多个）系统100可自动控制车辆200的制动、加速和/或转向。此外，系统100可分析所收集的数据，并基于所收集的数据的分析向驾驶者、乘客、用户或车辆（或其他车辆）的内部或外部的其他人发出警告、指示、建议、警报或指令。

[0091] 此外，与公开的实施例一致，由系统100提供的功能可使车辆200采取不同的动作，以在车道内以及/或者相对于其他车辆和/或对象导航车辆200。例如，系统100可调整车辆200相对于车辆200在其内行驶的车道和/或相对于被定位在车辆200附近的对象的定位，选择车辆200在行驶时使用的特定车道，以及响应于侵入的车辆（诸如试图移动到车辆200行驶的车道中的车辆）采取动作。另外，系统100可控制不同场景下（诸如当车辆200转弯时）的车辆200的速率。系统100可使车辆200模仿领先的车辆的动作或监测目标车辆并导航车辆200，使得车辆通过目标车辆。以下提供了与由系统100提供的各种实施例有关的额外细节。

[0092] 前视多成像系统

[0093] 如以上所讨论的，系统100可提供使用多相机系统的驾驶辅助功能。多相机系统可使用面向车辆的前向方向的一个或多个相机。在其他实施例中，多相机系统可包括面向车辆的侧面或车辆后方的一个或多个相机。在一个实施例中，例如，系统100可使用两相机成像系统，其中第一相机和第二相机（例如，图像捕获设备122和124）可被定位在车辆（例如，车辆200）的前面和/或侧面。第一相机可具有大于、小于第二相机的视场或与其部分重叠的视场。另外，第一相机可连接到第一图像处理器以对由第一相机提供的图像执行单目图像分析，并且第二相机可连接到第二图像处理器以对由第二相机提供的图像执行单目图像分析。第一图像处理器和第二图像处理器的输出（例如，被处理的信息）可组合。在一些实施例中，第二图像处理器可接收来自第一相机和第二相机二者的图像，以执行立体分析。在另一实施例中，系统100可使用其中相机中的每个具有不同的视场的三相机成像系统。因此，这样的系统可基于从位于车辆的前向和侧面二者的变化的距离处的对象得到的信息做出决策。对单目图像分析的参考可指基于从单一视点（例如，从单一相机）捕获的图像执行图像分析的实例。立体图像分析可指基于通过图像捕获参数的一个或多个变化捕获的两个或更多个图像来执行图像分析的实例。例如，适用于执行立体图像分析的所捕获的图像可包括使用不同焦距连同视差信息等从两个或更多个不同的定位捕获的图像、从不同的视场捕获的图像。

[0094] 例如，在一个实施例中，系统100可使用图像捕获设备122-126实施三相机配置。在这样的配置中，图像捕获设备122可提供窄视场（例如，34度或选自约20度到45度的范围的其他值等），图像捕获设备124可提供广视场（例如，150度或选自约100度到约180度的范围的其他值），以及图像捕获设备126可提供中间视场（例如，46度或选自约35度到约60度的范围的其他值）。在一些实施例中，图像捕获设备126可用作主相机或首要相机。图像捕获设备122-126可被定位在后视镜310之后，并且可被定位成基本上并排（例如，间隔6cm）。此外，在

IPR2025-00943
Tesla EX1002 Page 1107

一些实施例中，如以上所讨论的，一个或多个图像捕获设备122-126可被安装在与车辆200的挡风玻璃齐平的防眩板380之后。这样的防护可用于最小化从汽车内部的任何反射对图像捕获设备122-126的影响。

[0095] 在另一实施例中，如以上结合图3B和图3C所讨论的，广视场相机(例如，以上示例中的图像捕获设备124)可被安装为低于窄视场相机和主视场相机(例如，以上示例中的图像捕获设备122和126)。该配置可提供不受广视场相机影响的视线。为了降低反射，相机可被安装为接近车辆200的挡风玻璃，并且可在相机上包括偏振器以减弱反射光。

[0096] 三相机系统可提供某些性能特征。例如，一些实施例可包括基于来自另一相机的检测结果通过一个相机验证对象的检测的能力。在以上所讨论的三相机配置中，处理单元110可例如包括三个处理设备(例如，如以上所讨论的三个EyeQ系列的处理器芯片)，每个处理设备专用于处理由一个或多个图像捕获设备122-126捕获的图像。

[0097] 在三相机系统中，第一处理设备可接收来自主相机和窄视场相机二者的图像，并执行窄FOV相机的视觉处理，例如，以检测其他车辆、行人，车道标记、交通标志、交通灯和其他路对象。此外，第一处理设备可计算来自主相机的图像和来自窄相机的图像之间的像素的视差，以创建车辆200的环境的3D重建。随后，第一处理设备可将该3D重建与3D地图数据(例如，深度图)或与基于来自另外的相机的信息所计算的3D信息结合。

[0098] 第二处理设备可接收来自主相机的图像并执行视觉处理，以检测其他车辆、行人、车道标记、交通标志、交通灯和其他路对象。另外，第二处理设备可计算相机位移，并基于该位移计算连续的图像之间的像素的视差，并创建场景的3D重构(例如，运动恢复结构)。第二处理设备可将基于运动恢复结构的3D重构发送到第一处理设备，以与立体3D图像或与通过立体处理所获得的深度信息结合。

[0099] 第三处理设备可接收来自广FOV相机的图像并处理该图像，以检测车辆、行人、车道标记、交通标志、交通灯和其他路对象。第三处理设备可执行额外的处理指令，以分析图像从而辨认在图像中移动的对象，诸如改变车道的车辆、行人等。

[0100] 在一些实施例中，被单独捕获和处理的基于图像流的信息可提供用于提供系统中的冗余的机会。这样的冗余例如可包括使用第一图像捕获设备和由该设备处理的图像来验证和/或补充通过捕获和处理来自至少第二图像捕获设备的图像信息所获得信息。

[0101] 在一些实施例中，系统100可使用两个图像捕获设备(例如，图像捕获设备122和124)来为车辆200提供导航辅助，并且可使用第三图像捕获设备(例如，图像捕获设备126)来提供冗余并验证从其他两个图像捕获设备所接收的数据的分析。例如，在这样的配置中，图像捕获设备122和124可提供用于由系统100进行立体分析的图像以导航车辆200，而图像捕获设备126可提供用于由系统100进行单目分析的图像，以提供基于从图像捕获设备122和/或图像捕获设备124捕获的图像所获得的信息的冗余和验证。也就是说，图像捕获设备126(和相应的处理设备)可被认为提供了用于对从图像捕获设备122和124得到的分析进行检查的冗余子系统(例如，以提供自动应急制动(AEB)系统)。

[0102] 本领域中的技术人员将认识到，以上的相机配置、相机安置、相机的数量、相机位置等仅是示例。相对于整体系统所描述的这些组件和其他在不背离公开的实施例的范围的情况下可被装配并用在各种不同的配置中。涉及用于提供驾驶者辅助和/或自主车辆功能的多相机系统的使用的进一步细节如下。

IPR2025-00943
Tesla EX1002 Page 1108

[0103] 图4是存储器140和/或150的示例性功能框图，该存储器140和/或150能够储存/编程用于执行与公开的实施例一致的一个或多个操作的指令。虽然以下指的是存储器140，但本领域中的技术人员将认识到指令可被储存在存储器140和/或150中。

[0104] 如图4中所示，存储器140可储存单目图像分析模块402、立体图像分析模块404、速度和加速度模块406、以及导航响应模块408。公开的实施例不限于存储器140的任何特定的配置。此外，应用处理器180和/或图像处理器190可执行储存在被包括在存储器140中的模块402-408中任意一个中的指令。本领域中的技术人员将理解的是，在以下讨论中对处理单元110的引用可单独或统一指应用处理器180和图像处理器190。因此，以下过程中的任意步骤可由一个或多个处理设备执行。

[0105] 在一个实施例中，单目图像分析模块402可储存指令(诸如计算机视觉软件)，该指令在由处理单元110执行时，对由图像捕获设备122、124和126中的一个所采集的一组图像进行单目图像分析。在一些实施例中，处理单元110可使来自一组图像的信息与附加的感觉信息(例如，来自雷达的信息)结合，以执行单目图像分析。如以下结合图5A-图5D所描述的，单目图像分析模块402可包括用于检测该组图像内的一组特征(诸如车道标记、车辆、行人、路标志、高速公路出口匝道、交通灯、危险对象和与车辆的环境相关联的任何其他特征)的指令。在一些实施例中，对于任何给定的帧，车辆的环境可包括由成像设备中的任意一个或由其任意组合覆盖的区域或空间(成像设备的FOV)。如以下结合导航响应模块408所讨论的，基于该分析，系统100(例如，经由处理单元110)可在车辆200中引起一个或多个导航响应，诸如转弯、车道移位、加速度的变化等。此外，基于该分析，系统100(例如，经由处理单元110)可引起一个或多个导航和/或车辆控制响应，包括但不限于控制车辆200以及以上所讨论的任何相关联的系统和/或子系统的各种方式。在一些实施例中，以上所提到的响应或动作中任意一个或其全部为了简单起见可被称作导航动作或导航响应。

[0106] 在一个实施例中，立体图像分析模块404可储存指令(诸如计算机视觉软件)，该指令在由处理单元110执行时，对由选自图像捕获设备122、124和126中的任意一个的图像捕获设备的组合所采集的第一组图像和第二组图像执行立体图像分析。在一些实施例中，处理单元110可使来自第一组图像和第二组图像的信息与附加的感觉信息(例如，来自雷达的信息)结合，以执行立体图像分析。例如，立体图像分析模块404可包括用于基于由图像捕获设备124所采集的第一组图像和由图像捕获设备126所采集的第二组图像执行立体图像分析的指令。如以下结合图6所描述的，立体图像分析模块404可包括用于检测第一组图像和第二组图像内的一组特征(诸如车道标记、车辆、行人、路标志、高速公路出口匝道、交通灯、危险对象等)的指令。如以下结合导航响应模块408所讨论的，基于该分析，处理单元110可在车辆200中引起一个或多个导航响应，诸如转弯、车道移位、加速度的变化等。

[0107] 在一个实施例中，速度和加速度模块406可储存被配置为对从车辆200中的被配置为引起车辆200的速度和/或加速度的改变的一个或多个计算设备和机电设备所接收的数据进行分析。例如，处理单元110可执行与速度和加速度模块406相关联的指令，以基于通过单目图像分析模块402和/或立体图像分析模块404的执行得出的数据计算车辆200的目标速率。这样的数据例如可包括目标定位、速度和/或加速度、车辆200相对于附近的车辆、行人或路对象的定位和/或速率、车辆200相对于路的车道标记等的定位信息。另外，处理单元110可基于感觉输入(例如，来自雷达的信息)和来自车辆200的其他系统(诸如车辆200的节

IPR2025-00943
Tesla EX1002 Page 1109

流系统220、制动系统230和/或转向系统240)的输入计算车辆200的目标速率。基于所计算的目标速率，处理单元110可将电信号传输到车辆200的节流系统220、制动系统230和/或转向系统240，以例如通过物理压下制动器或放松车辆200的加速器来触发速度和/或加速度的改变。

[0108] 在一个实施例中，导航响应模块408可储存可由处理单元110执行的软件，以基于通过单目图像分析模块402和/或立体图像分析模块404的执行得出的数据确定所需的导航响应。这样的数据可包括与附近的车辆、行人和路对象相关联的定位信息和速率信息、车辆200的目标定位信息等。另外，在一些实施例中，导航响应可(部分或全部)基于地图数据、车辆200的预定定位、以及/或者车辆200与通过单目图像分析模块402和/或立体图像分析模块404的执行所检测到的一个或多个对象之间的相对速度或相对加速度。导航响应模块408也可基于感觉输入(例如，来自雷达的信息)和来自车辆200的其他系统(诸如车辆200的节流系统220、制动系统230和转向系统240)的输入确定所需的导航响应。基于所需的导航响应，处理单元110可将电信号传输到车辆200的节流系统220、制动系统230和转向系统240以例如通过转动车辆200的方向盘来实现预定角度的旋转从而触发所需的导航响应。在一些实施例中，处理单元110可使用导航响应模块408的输出(例如，所需的导航响应)作为速度和加速度模块406的执行的输入，以用于计算车辆200的速度的改变。

[0109] 图5A是示出了与公开的实施例一致的用于基于单目图像分析引起一个或多个导航响应的示例性过程500A的流程图。在步骤510，处理单元110可经由处理单元110和图像采集单元120之间的数据接口128接收多个图像。例如，被包括在图像采集单元120中的相机(诸如具有视场202的图像捕获设备122)可捕获车辆200的前向区域(或者例如，车辆的侧面或其后方)的多个图像，并通过(例如，数字、有线、USB、无线、蓝牙等)数据连接将它们传输到处理单元110。如以下结合图5B-图5D进一步详细描述的，在步骤520，处理单元110可使单目图像分析模块402执行对多个图像的分析。通过执行该分析，处理单元110可检测该组图像内的一组特征，诸如车道标记、车辆、行人、路标志、高速公路出口匝道、交通灯等。

[0110] 在步骤520，处理单元110也可使单目图像分析模块402执行对各种路危险物(卡车轮胎的部分、倒下的路标志、松散的货物、小动物等)进行检测。路危险物可能在结构、形状、大小和颜色上有所不同，这可能使对这样的危险物做出决策更具挑战性。在一些实施例中，处理单元110可使单目图像分析模块402对多个图像执行多帧分析，以检测路危险物。例如，处理单元110可估计连续的图像帧之间的相机运动，并计算帧之间的像素的视差，以构建路或路的某些区域(例如，在图像数据周围所界定的疑似有路危险物的兴趣区域内)的3D地图。随后，处理单元110可使用该3D地图来检测路危险物。例如，分割技术和分类技术可用于检测路表面以及存在于路表面之上的危险物。例如，如以下另外详细解释的，系统可检测路表面的特征，诸如画的车道线或车道标记，并且另外可检测存在于路表面以上的道路的特征，诸如施工区中的圆锥体/圆筒体/杆/障碍物、护栏、其他车辆、路标志等。

[0111] 在步骤530，处理单元110可使导航响应模块408执行基于在步骤520进行的分析和如以上结合图4所描述的技术在车辆200中引起一个或多个导航响应。例如，导航响应可包括转弯、车道移位、加速度的改变等。在一些实施例中，处理单元110可使用通过速度和加速度模块406的执行得出的数据，以引起一个或多个导航响应。另外，多个导航响应可同时、依次或它们的任意组合地发生。例如，处理单元110可使车辆200移过一个车道，并随后例如通

IPR2025-00943
Tesla EX1002 Page 1110

过依次将控制信号传输到车辆200的转向系统240和节流系统220来进行加速。可替代地，处理单元110可使车辆200制动，同时例如通过同时将控制信号传输到车辆200的制动系统230和转向系统240移位车道。

[0112] 图5B是示出了与公开的实施例一致的用于在一组图像中检测一个或多个车辆和/或行人的示例性过程500B的流程图。处理单元110可使单目图像分析模块402执行过程500B。在步骤540，处理单元110可确定表示可能的车辆和/或行人的一组候选对象。例如，处理单元110可扫描一个或多个图像，将图像与一个或多个预定图案进行比较，并在可能包含感兴趣的对象(例如，车辆、行人或它们的部分)的可能的位置内辨认每个图像。预定的图案可通过这样的方式来设计，以实现高比率的“错误命中”和低比率的“遗漏”。例如，处理单元110可使用类似于预定的图案的低阈值，以用于辨认候选对象为可能的车辆或行人。这样做可使处理单元110能够降低遗漏(例如，没有辨认出)表示车辆或行人的候选对象的可能性。

[0113] 在步骤542，处理单元110可基于分类标准过滤该组候选对象，以排除某些候选(例如，不相关的或不太相关的对象)。这样的标准可从与储存在数据库(例如，储存在存储器140中的数据库)中的对象类型相关联的各种属性中得出。属性可包括对象的形状、尺寸、纹理、(例如，相对于车辆200的)定位等。因此，处理单元110可使用一组或多组标准，以排斥该组候选对象中的错误的候选。

[0114] 在步骤544，处理单元110可分析多帧图像，以确定在该组候选对象中的对象是否表示车辆和/或行人。例如，处理单元110可跟踪在连续的帧中被检测的候选对象，并累积与被检测的对象相关联的逐帧数据(例如，相对于车辆200的大小、定位等)。另外，处理单元110可估计被检测的对象的参数，并将对象的逐帧定位数据与预测的定位进行比较。

[0115] 在步骤546，处理单元110可构建关于被检测的对象的一组测量值。这样的测量值例如可包括与被检测的对象相关联的(相对于车辆200)的大小、定位、速度和加速度的值。在一些实施例中，处理单元110可基于使用一系列基于时间的观察的估计技术(诸如，卡尔曼滤波器或线性二次估计(LQE))，和/或基于不同对象类型(例如，汽车、卡车、行人、自行车、路标志等)的可用建模数据来构建测量值。卡尔曼滤波器可基于对象的比例的测量值，其中该比例测量值与碰撞时间(例如，车辆200到达对象的时间量)成比例。因此，通过执行步骤540-546，处理单元110可辨认出现在该组捕获的图像内的车辆和行人，并得出与该车辆和行人相关联的信息(例如，定位、速率、大小)。如以上结合图5A所描述的，基于辨认和得出的信息，处理单元110可在车辆200中引起一个或多个导航响应。

[0116] 在步骤548，处理单元110可执行一个或多个图像的光流分析，以降低检测“错误命中”和遗漏表示车辆或行人的候选对象的可能性。例如，光流分析可指分析相对于与其他车辆和行人相关联的一个或多个图像中的车辆200的运动图案，并且该运动图案与路表面的运动不同。处理单元110可通过观察在不同时间被捕获的多个图像帧中的对象的不同定位来计算候选对象的运动。处理单元110可使用定位和时间的值作为用于计算候选对象的运动的数学模型的输入。因此，光流分析可提供检测车辆200附近的车辆和行人的另一种方法。处理单元110可结合步骤540-546执行光流分析，以提供用于检测车辆和行人的冗余并提高系统100的可靠性。

[0117] 图5C是示出了与公开的实施例一致的用于在一组图像中检测路标记和/或车道几何形状信息的示例性过程500C的流程图。处理单元110可执行单目图像分析模块402以实施

IPR2025-00943
Tesla EX1002 Page 1111

过程500C。在步骤550，处理单元110可通过扫描一个或多个图像检测一组对象。为了检测车道标记、车道几何形状信息和其他有关的路标记的节段，处理单元110可过滤该组对象，以排除确定是不相关的那些对象(例如，小坑、小石头等)。在步骤552，处理单元110可将在步骤550中检测到的属于相同的路标记或车道标记的节段组合在一起。基于该分组，处理单元110可开发表示被检测到的节段的模型，诸如数学模型。

[0118] 在步骤554，处理单元110可构建与被检测到的节段相关联的一组测量值。在一些实施例中，处理单元110可将从图像平面检测到的节段的投影创建到实际平面上。投影可使用具有对应于物理属性(诸如被检测的路的定位、斜率、曲率和曲率导数)的系数的三次多项式来表征。在生成投影时，处理单元110可考虑路表面的改变以及与车辆200相关联的俯仰和滚转的速率。另外，处理单元110可通过分析路表面上现有的定位和运动的线索建模路标高。此外，处理单元110可通过跟踪一个或多个图像中的一组特征点估计与车辆200相关联的俯仰和滚转的速率。

[0119] 在步骤556，处理单元110可例如通过跟踪在连续的图像帧中的被检测到的节段并累积与被检测到的节段相关联的逐帧数据来执行多帧分析。在处理单元110执行多帧分析时，在步骤554被构建的该组测量值可变得更加可靠并与越来越高的置信水平相关联。因此，通过执行步骤550-556，处理单元110可辨认在该组被捕获的图像内出现的路标记，并得出车道几何形状的信息。如以上结合图5A所描述的，基于辨认和得出的信息，处理单元110可在车辆200中引起一个或多个导航响应。

[0120] 在步骤558，处理单元110可考虑额外的信息源，以进一步开发用于车辆200的周围环境的背景中的车辆200的安全模型。处理单元110可使用安全模型来界定其中系统100能够以安全模式执行车辆200的自主控制的背景。为了开发安全模型，在一些实施例中，处理单元110可考虑其他车辆的定位和运动、被检测到的路边缘和障碍物、和/或从地图数据(诸如来自地图数据库160的数据)提取的大致的路形状的描述。通过考虑额外的信息源，处理单元110可提供用于检测路标记和车道几何形状的冗余并增加了系统100的可靠性。

[0121] 图5D是示出了与公开的实施例一致的用于在一组图像中检测交通灯的示例性过程500D的流程图。处理单元110可使单目图像分析模块402执行过程500D。在步骤560，处理单元110可扫描该组图像并辨认在可能包含交通灯的图像中的位置处出现的对象。例如，处理单元110可过滤被辨认的对象，以构建排除了那些不可能对应于交通灯的对象的一组候选对象。过滤可基于与交通灯相关联的各种属性(诸如形状、尺寸、纹理、(相对于车辆200的)定位等)来完成。这样的属性可基于交通灯和交通控制信号的多个示例，并且可被储存在数据库中。在一些实施例中，处理单元110可对反映了可能的交通灯的该组候选对象执行多帧分析。例如，处理单元110可跟踪在连续的图像帧中的候选对象，估计候选对象的真实定位，并过滤出移动的那些对象(其不可能是交通灯)。在一些实施例中，处理单元110可对候选对象执行颜色分析，并辨认在可能的交通灯内出现的被检测的颜色的相对定位。

[0122] 在步骤562，处理单元110可分析交叉口的几何形状。该分析可基于以下项的任意组合：(i)在车辆200的任意侧面上所检测到的车道的数量，(ii)路上所检测到的标记(诸如箭头标记)，以及(iii)从地图数据(诸如来自地图数据库160的数据)提取的交叉口的描述。处理单元110可使用通过执行单目分析模块402得到的信息来进行该分析。另外，处理单元110可确定在步骤560检测到的交通灯和在车辆200附近出现的车道之间的对应关系。

IPR2025-00943
Tesla EX1002 Page 1112

[0123] 在步骤564，当车辆200接近交叉口时，处理单元110可更新与所分析的交叉口几何形状和被检测到的交通灯相关联的置信水平。例如，相较于在交叉口实际出现的交通灯的数量，估计在交叉口出现的交通灯的数量可能影响置信水平。因此，基于置信水平，处理单元110可将控制委托给车辆200的驾驶者，以便改进安全条件。通过执行步骤560-564，处理单元110可辨认在该组被捕获的图像内出现的交通灯，并分析交叉口的几何形状的信息。如以上结合图5A所描述的，基于辨认和分析，处理单元110可在车辆200中引起一个或多个导航响应。

[0124] 图5E是示出了与公开的实施例一致的用于基于车辆路径在车辆200中引起一个或多个导航响应的示例性过程500E的流程图。在步骤570，处理单元110可构建与车辆200相关联的初始的车辆路径。车辆路径可使用以坐标(x,z)表达的一组点来表示，以及在一些实施例中，该组点中的两点之间的距离d_i可落在1米到5米的范围内，或者在其他实施例中，定时测量可用于提供该组点，其中该组点中的点彼此间隔0.5秒到5秒(例如，间隔1.5秒)。在一个实施例中，处理单元110可使用两个多项式(polynomials)(诸如左侧路和右侧路的多项式)来构建初始的车辆路径。如果有的话(零偏移可对应于车道中间的行进)，处理单元110可计算两个多项式之间的几何中点，并使被包括在得到的车辆路径中的每个点偏移预定的偏移(例如，智能车道偏移)。偏移可以在垂直于车辆路径中的任意两点之间的节段的方向上。在另一实施例中，处理单元110可使用一个多项式和所估计的车道宽度，以使车辆路径的每个点偏移所估计的车道宽度的一半加预定的偏移(例如，智能车道偏移)。

[0125] 在步骤572，处理单元110可更新在步骤570所构建的车辆路径。处理单元110可使用更高的分辨率重新构建在步骤570所构建的车辆路径，使得表示车辆路径的该组点中的两点之间的距离d_k小于以上所描述的距离d_i。例如，距离d_k可落在0.1米到0.3米的范围内。处理单元110可使用抛物线样条算法来重新构建车辆路径，这可得出对应于车辆路径的总长度(即，基于表示车辆路径的该组点)的累积的距离向量S。

[0126] 在步骤574，处理单元110可基于在步骤572所构建的更新的车辆路径确定(以坐标表达为(x1,z1)的)先行点(look-ahead point)。处理单元110可从累积的距离向量S提取先行点，并且该先行点可与先行距离和先行时间相关联。可具有10米到20米范围的下限的先行距离可作为车辆200的速率和先行时间的乘积来计算。例如，随着车辆200的速度降低，先行距离也可能减少(例如，直至该先行距离达到下限)。可介于0.5秒到1.5秒范围内的先行时间可与一个或多个控制回路(诸如航向误差跟踪控制回路)的增益成反比，该一个或多个控制回路与在车辆200中引起导航响应相关联。例如，航向误差跟踪控制回路的增益可取决于汽车横向动力、转向执行器回路、偏航速率回路的带宽等。因此，航向误差跟踪控制回路的增益越高，先行时间就越低。

[0127] 在步骤576，处理单元110可基于在步骤574确定的先行点确定航向误差和偏航速率命令。处理单元110可通过计算先行点的反正切(例如，arctan(x1/z1))确定航向误差。处理单元110可确定偏航速率命令为航向误差和高水平控制增益的乘积。高水平控制增益可等于：(2/先行时间)，如果先行距离不在下限上。否则，高水平控制增益可等于：(2*车辆200的速率/先行距离)。

[0128] 图5F是示出了与公开的实施例一致的用于确定领先的车辆是否正改变车道的示例性过程500F的流程图。在步骤580，处理单元110可确定与领先的车辆(例如，在车辆200的

IPR2025-00943
Tesla EX1002 Page 1113

前方行进的车辆)相关联的导航信息。例如,处理单元110可使用以上结合图5A和图5B所描述的技术,确定领先的车辆的定位、速度(例如,方向和速率)和/或加速度。处理单元110也可使用以上结合图5E所描述的技术,确定一个或多个路多项式、(与车辆200相关联的)先行点和/或蜗行踪迹(snail trail)(例如,描述领先的车辆所采取的路径的一组点)。

[0129] 在步骤582,处理单元110可分析在步骤580所确定的导航信息。在一个实施例中,处理单元110可计算蜗行踪迹和(例如,沿着该踪迹的)路多项式之间的距离。如果沿着踪迹的该距离的方差超过预定的阈值(例如,直路上0.1米到0.2米、适度弯曲的路上0.3米到0.4米、以及具有尖锐曲线的路上0.5米到0.6米),则处理单元110可确定领先的车辆可能正改变车道。在检测到多个车辆在车辆200的前方行进的情况下,处理单元110可比较与每个车辆相关联的蜗行踪迹。基于该比较,处理单元110可确定其蜗行踪迹与其他车辆的蜗行踪迹不匹配的车辆可能正改变车道。另外,处理单元110可将(与领先的车辆相关联的)蜗行踪迹的曲率与领先的车辆行进的路段的预期的曲率进行比较。预期的曲率可从地图数据(例如,来自地图数据库160的数据)、路多项式、其他车辆的蜗行踪迹、关于路的先验知识等提取出来。如果蜗行踪迹和曲率和路段的预期的曲率的差值超过预定的阈值,则处理单元110可确定领先的车辆可能正改变车道。

[0130] 在另一实施例中,处理单元110可在指定的时间段(例如,0.5秒到1.5秒)内将领先的车辆的瞬时定位与(与车辆200相关联的)先行点进行比较。如果在指定的时间段期间,领先的车辆的瞬时定位和先行点之间的距离变化,并且累积的变化的和超过预定的阈值(例如,直路上0.3米到0.4米、适度弯曲的路上0.7米到0.8米、以及具有尖锐曲线的路上1.3米到1.7米),则处理单元110可确定领先的车辆可能正改变车道。在另一实施例中,处理单元110可通过将沿着踪迹行进的横向距离与蜗行踪迹的预期的曲率进行比较来分析该蜗行踪迹的几何形状。期望的曲率半径可根据以下计算来确定:$(\delta_z^2+\delta_x^2)/2/(\delta_x)$,其中,$\delta_x$表示行进的横向距离,以及$\delta_z$表示行进的纵向距离。如果行进的横向距离和期望的曲率之间的差值超过预定的阈值(例如,500米到700米),则处理单元110可确定领先的车辆可能正改变车道。在另一实施例中,处理单元110可分析领先的车辆的定位。如果领先的车辆的定位掩盖了道路多项式(例如,领先的车辆重叠在路多项式的顶部上),则处理单元110可确定领先的车辆可能正改变车道。在领先的车辆的定位的状况是检测到另一车辆在该领先的车辆的前方且两个车辆的蜗行踪迹不平行的情况下,处理单元110可确定(较近的)领先的车辆可能正改变车道。

[0131] 在步骤584,处理单元110可基于在步骤582所执行的分析确定领先的车辆200是否正改变车道。例如,处理单元110可基于在步骤582执行的各个分析的加权平均值做出确定。在这样的方案下,例如,处理单元110基于特定类型的分析确定领先的车辆可能正改变车道的决定可被分配值"1"(以及"0"表示领先的车辆不像正改变车道的确定)。在步骤582执行的不同分析可被分配不同的权重,并且公开的实施例不限于分析和权重的任何特定组合。

[0132] 图6是示出了与公开的实施例一致的用于基于立体图像分析引起一个或多个导航响应的示例性过程600的流程图。在步骤610,处理单元110可经由数据接口128接收第一多个图像和第二多个图像。例如,被包括在图像采集单元120中的相机(诸如具有视场202和204的图像捕获设备122和124)可捕获车辆200的前向区域的第一多个图像和第二多个图像,并通过数字连接(例如,USB、无线、蓝牙等)将它们传输给处理单元110。在一些实施例

IPR2025-00943
Tesla EX1002 Page 1114

中，处理单元110可经由两个或更多个数据接口接收第一多个图像和第二多个图像。公开的实施例不限于任何特定的数据接口配置或协议。

[0133] 在步骤620，处理单元110可使立体图像分析模块404对第一多个图像和第二多个图像执行立体图像分析，以创建在车辆的前面的路的3D地图并检测图像内的特征，诸如车道标记、车辆、行人、路标志、高速公路出口匝道、交通灯、路危险物等。立体图像分析能够以类似于以上结合图5A-图5D所描述的步骤的方式来执行。例如，处理单元110可使立体图像分析模块404执行对第一多个图像和第二多个图像内的候选对象(例如，车辆、行人、路标记、交通灯、路危险物等)进行检测，基于各种标准过滤出候选对象的子集，以及执行多帧分析、构建测量值、和确定关于剩余的候选对象的置信水平。在执行以上步骤时，处理单元110可考虑来自第一多个图像和第二多个图像二者的信息，而不是单独来自一组图像的信息。例如，处理单元110可对出现在第一多个图像和第二多个图像二者中的候选对象分析像素级数据(或来自两个捕获的图像流中的其他数据子集)上的差异。另一示例是，处理单元110可通过观察到对象出现在多个图像中的一个中但未出现在其他图像中或者对象相对于其他的可能相对于出现在两个图像流中的对象存在的差异来估计(例如，相对于车辆200的)候选对象的定位和/或速度。例如，相对于车辆200的定位、速度和/或加速度可基于与出现在一个或两个图像流中的对象相关联的特征的轨迹、定位、移动特性等来确定。

[0134] 在步骤630，处理单元110可使导航响应模块408执行基于在步骤620进行的分析和如以上结合图4所描述的技术在车辆200中引起一个或多个导航响应。例如，导航响应可包括转弯、车道移位、加速度的改变、速度的改变、制动等。在一些实施例中，处理单元110可使用通过速度和加速度模块406的执行得出的数据来引起一个或多个导航响应。另外，多个导航响应可同时、依次或它们的任意组合地发生。

[0135] 图7是示出了与公开的实施例一致的用于基于三组图像的分析引起一个或多个导航响应的示例性过程700的流程图。在步骤710，处理单元110可经由数据接口128接收第一多个图像、第二多个图像和第三多个图像。例如，被包括在图像采集单元120中的相机(诸如具有视场202、204和206的图像捕获设备122、124和126)可捕获车辆200的前向区域和/或侧面的第一多个图像、第二多个图像和第三多个图像，并通过数字连接(例如，USB、无线、蓝牙等)将它们传输给处理单元110。在一些实施例中，处理单元110可经由三个或更多个数据接口接收第一多个图像、第二多个图像和第三多个图像。例如，图像捕获设备122、124和126中的每个可具有用于将数据传递给处理单元110的相关数据接口。公开的实施例不限于任何特定的数据接口配置或协议。

[0136] 在步骤720，处理单元110可分析第一多个图像、第二多个图像和第三多个图像，以检测图像内的特征，诸如车道标记、车辆、行人、路标志、高速公路出口匝道、交通灯、路危险物等。该分析可以以类似于以上结合图5A-图5D及图6所描述的步骤的方式来执行。例如，处理单元110可对第一多个图像、第二多个图像和第三多个图像的每个(例如，通过执行单目图像分析模块402并基于以上结合图5A-图5D所描述的步骤)执行单目图像分析。可替代地，处理单元110可对第一多个图像和第二多个图像、第二多个图像和第三多个图像、和/或第一多个图像和第三多个图像(例如，通过执行单目图像分析模块404并基于以上结合图6所描述的步骤)执行立体图像分析。对应于第一多个图像、第二多个图像和/或第三多个图像的分析所处理的信息可组合。在一些实施例中，处理单元110可执行单目图像分析和立体图

IPR2025-00943
Tesla EX1002 Page 1115

像分析的组合。例如，处理单元110可对第一多个图像(例如，通过执行单目图像分析模块402)执行单目图像分析，并对第二多个图像和第三多个图像(例如，通过执行立体图像分析模块404)执行立体图像分析。图像捕获设备122、124和126的配置-包括它们各自的位置和视场202、204和206-可影响对第一多个图像、第二多个图像和第三多个图像所进行的分析的类型。公开的实施例不限于图像捕获设备122、124和126的特定配置或在第一多个图像、第二多个图像和第三多个图像上所进行的分析的类型。

[0137] 在一些实施例中，处理单元110可在系统100上基于在步骤710和720所采集并分析的图像执行测试。这样的测试可提供关于图像捕获设备122、124和126的某些配置的系统100的整体性能的指示物。例如，处理单元110可确定“错误命中”(例如，系统100没有正确确定车辆或行人的存在的情况)和“遗漏”的比例。

[0138] 在步骤730，处理单元110可基于从第一多个图像、第二多个图像和第三多个图像中的两者得出的信息在车辆200中引起一个或多个导航响应。第一多个图像、第二多个图像和第三多个图像中的两者的选择可取决于各种因素，诸如例如，在多个图像的每个中所检测到的对象的数量、类型和大小。处理单元110也可基于图像的质量和分辨率、图像中反映的有效视场、被捕获的帧的数量、一个或多个感兴趣对象实际出现在帧中的程度(例如，对象出现在其中的帧的百分比、出现在每个这样的帧中的对象的比例等)等做出选择。

[0139] 在一些实施例中，处理单元110可通过确定从一个图像源得出的信息与从其他图像源得出的信息一致的程度，选择从第一多个图像、第二多个图像和第三多个图像中的两者得出的信息。例如，处理单元110可(通过单目分析、立体分析或二者的任意组合)组合从图像捕获设备122、124和126中的每个得出的被处理的信息，并确定在从图像捕获设备122、124和126中的每个所捕获的图像中一致的视觉指示物(例如，车道标记、被检测到的车辆以及其位置和/或路径、被检测到的交通灯等)。处理单元110也可排除在被捕获的图像中不一致的信息(例如，改变车道的车辆、指示车辆太接近车辆200的车道模型等)。因此，处理单元110可基于一致的确定和不一致的信息选择从第一多个图像、第二多个图像和第三多个图像中的两者得出的信息。

[0140] 例如，导航响应可包括转弯、车道移位、加速度的改变等。处理单元110可基于在步骤720所执行的分析和以上结合图4所描述的技术引起一个或多个导航响应。处理单元110也可使用通过执行速度和加速度模块406得出的数据来引起一个或多个导航响应。在一些实施例中，处理单元110可基于车辆200和第一多个图像、第二多个图像以及第三多个图像中任意一个内所检测到的对象之间的相对定位、相对速度和/或相对加速度来引起一个或多个导航响应。多个导航响应可同时、依次或它们的任意组合地发生。

[0141] 车道限制预测

[0142] 系统100可给车辆200提供车道限制预测技术，使得系统预测一个或多个车道限制的位置。

[0143] 在一些实施例中，车道限制预测可被理解为用于预测车道限制的位置的一种或多种技术。在一些实施例中，车道限制可被界定为车道、道路、停车位、或者能够使车辆行进的其他路径或空间的边缘处的分界线。在一些情况下，车道限制可被标记在道路上，诸如通过画的线、物理障碍物、博茨点或者一个或多个额外的指示物。在一些实施例中，车道限制可由道路、路边、护栏等的边缘界定。在一些实施例中，车道限制可由一个或多个其他车辆(诸

IPR2025-00943
Tesla EX1002 Page 1116

如临近车道成排平行停放的车辆，或停车场中成排停放的车辆）界定。在一些实施例中，车道限制可不做标记，诸如在未标记的道路上或在没有车道线的停车场中。

[0144] 在一些实施例中，自驾驶系统或驾驶者辅助系统可使用图像捕获设备和图像分析，以基于在图像中所检测的对象（诸如车道标记、道路的边缘和/或其他车辆）确定一个或多个车道限制的位置。在大多驾驶情况下，系统可以能够基于所捕获的图像的可见的特性（例如，车道线等）检测至少两个车道限制——车辆行进的车道的每个侧面一个。

[0145] 然而，在某些情况下，包括在车道的一侧上或车辆200的一侧上的仅一个车道限制可基于所捕获的图像中的可见的特性（例如，车道线等）被充分检测到的那些情况，系统100可预测在车道的另一侧上或车辆的另一侧上的车道限制的定位，使车辆能够在由所检测到的限制和与其相对的所预测的限制界定的车道内继续行进。

[0146] 在一些实施例中，车道限制预测可由以满足一个或多个条件为前提的系统100执行。例如，在一些实施例中，当先前所检测的车道限制随着车辆200沿着道路前进被停止检测时，系统100预测车道限制的位置。因此，在一些实施例中，如果车道标记或其他所检测到的车道限制突然结束或变成不可检测，则可执行车道限制预测。在其他示例中，当系统确定护栏阴影条件存在时，系统100预测车道限制的位置，使得护栏存在于道路的附近或者护栏在道路上投射阴影。

[0147] 在一些实施例中，当系统100具有对关于车辆200行进的车道宽度的有效且足够的历史信息访问时，系统100可仅执行车道限制预测；在车辆200的一侧上的车道限制随着车辆200沿着道路前进被新检测到的情况下，在一些实施例中，系统100可不执行车道限制预测。

[0148] 在一些实施例中，车道限制预测可在最大行进距离内被执行，使得如果车辆200在系统100连续执行车道限制预测时行进了比预定的最大行进距离更大的距离，则车道限制预测随后可停止。在一些实施例中，预测的最大行进距离可根据由系统100检测的一个或多个各种因素来设置或修改。在一些实施例中，预测的最大行进距离可根据由系统100检测的一个或多个各种因素而折向（例如，其可在功能上被设置为“无限”）。

[0149] 在一些实施例中，车道限制预测可根据由系统100检测的各种取消条件而停止。例如，在一些实施例中，如果车辆200的驾驶者改变车道，则车道限制预测可停止。在一些实施例中，如果车辆200在道路中横穿超过预定的曲线锐度阈值的曲线，则车道限制预测可停止。

[0150] 图8A-图8C图示了车辆200在道路800上行进的实施例。在一些实施例中，车辆200可具有由一个或多个图像捕获设备或者相机（诸如图像捕获设备122、124和/或126）提供的一个或多个视场。在所描绘的示例中，如在车辆200的前面形成向上开口的三角形的成角度的线所描绘的，车辆200可具有从车辆朝前面向的视场；在一些实施例中，车辆200可具有从车辆200面向多个不同的方向的一个或多个视场。在所描绘的示例中，可使用用于车道限制预测（例如，当没有检测到车道限制时，预测车道限制的位置）和在由预测的车道限制界定的车道内操作车辆200的公开的系统和方法。如图所示，道路800可具有车辆200行进的车道802。车道802可在第一侧面上由第一车道限制804界定，以及在第一侧面相对的第二侧面上由车道限制806界定。如示出车道限制804和806的粗体线所指示的，这些车道限制804和806可由系统100检测。车辆200可在车辆200与第一车道限制804相同的侧面上具有第一车辆侧

IPR2025-00943
Tesla EX1002 Page 1117

面808，以及在车辆200与第二车道限制806相同的侧面上具有第二车辆侧面810。

[0151] 处理单元110可被配置为基于由图像捕获设备122-126采集的多个图像确定第一车道限制804和第二车道限制806，该多个图像可经由数据接口128由处理单元110接收。根据一些实施例，第一车道限制804和/或第二车道限制806可由可见的车道分界线(诸如像车道线803a(实线)和803b(虚线)的路表面上标记的虚线或实线)辨认。另外或可替代地，第一车道限制804和/或第二车道限制806可由障碍物或路表面的边缘辨认。另外或可替代地，第一车道限制804和/或第二车道限制806可标记(例如，博茨点)辨认。根据一些实施例，处理单元110可通过辨认路表面宽度的中点确定第一车道限制804和/或第二车道限制806。根据一些实施例，如果处理单元110仅辨认第一车道限制804，则处理单元110可估计第二车道限制806，诸如基于所估计的车道宽度或路宽度。例如，当指定路车道的线没有被画出或通过其他方式标记出来时，处理单元110可通过这种方式预测车道限制。

[0152] 在图8A-图8C中，车道限制804和806分别通过实线和虚线803a和803b被标记在道路800上。在朝向图8A-图8C的中心的道路800的部分中，对应于车道限制806的虚线803b没有被画在道路800上。在图8A-图8C中，车道限制804和806通过系统100的检测由覆盖在车辆200的前面的相应的车道线的粗体线描绘。如图8A中的粗体线所描绘的，系统100在车辆200的前面的一定可见的距离内检测作为车辆200的左侧的实线车道限制的车道限制804。类似地，系统100在车辆200的前面的一定可见的距离内检测作为车辆200的右侧的虚线车道限制的车道限制806。根据车道限制804和806的检测，系统100使车辆200在由车道限制804和806界定的车道802内行进。

[0153] 图8B描绘了在图8A中所描绘的时间点之后在车道802中沿着道路800继续前进的车辆200。当车辆200到达虚线803b没有被画在道路800上的道路800的部分时，车道限制预测可由系统100执行。如图所示，系统100预测车道限制806的位置位于距相对的车道限制804的距离812处。在一些实施例中，该距离812可基于所检测到的历史信息以及由系统100储存的关于车道802的先前宽度和/或车道大致上的平均历史宽度来确定。如图所示，系统100基于在车道限制预测被激活之前车道限制806被系统100检测到的最后一次由虚的车道线803b标记车道限制806的事实来预测车道限制806继续是虚的车道限制。根据车道限制804的继续的检测和车道限制806的预测，系统100使车辆200在由车道限制804和806界定的车道802内继续行进。在一些实施例中，当车道限制预测被激活时，并且可能在某些情况(例如，附近有其他车辆)下，某默认类型的车道限制用于控制车辆的运动，该默认类型的车道限制不一定与在激活车道限制预测之前所检测到的最后的车道限制相同。例如，在一些实施例中，系统可默认为预测实线车道限制而不是虚线车道限制。

[0154] 图8C描绘了在图8B中所描绘的时间点之后在车道802中沿着道路800继续前进的车辆200。当车辆200到达虚线803b又再次被画在道路800上的道路800的部分时，由于车道限制806作为虚的车道限制又再次由系统100检测到(而不是被预测)，因此车道限制预测可由系统100停止执行。如图8C中所描绘的，车辆200在执行车道限制预测时行进了距离814。在一些实施例中，车道限制预测仅可在预定的最大距离(和/或预定的最大时间)内被连续或基本上连续地执行。在图8C所描绘的示例中，距离814小于用于执行车道限制预测的预定的最大距离，使得车道限制预测随着车辆200横穿其在图8B中的定位和其在图8C中的定位之间的距离能够被连续执行。在一些实施例中，用于执行车道限制预测的预定的最大距离

IPR2025-00943
Tesla EX1002 Page 1118

可被设置为18米的默认值。在一些实施例中，预定的最大距离可由车辆制造商、驾驶者和/或任何其他人员改变。

[0155] 图9A-图9C图示了车辆200在道路900上行进的实施例。在所描绘的示例中，可使用用于车道限制预测(例如，当没有检测到车道限制时，预测车道限制的位置)和在由预测的车道限制界定的车道内操作车辆200的公开的系统和方法。图9A-图9C中所描绘的示例不同于图8A-图8C中所描绘的示例至少在于在道路900上两个车道合并成一个车道。如以下更详细解释的，在一些实施例中，当两个车道合并成一个车道且车道线在合并第二车道的车道的侧面上未被检测到时，可执行车道限制预测。

[0156] 如图所示，道路900可具有车辆200行进的车道902。车道902可在第一侧面上由第一车道限制904界定，以及在第一侧面相对的第二侧面上由车道限制906界定。如示出车道限制904和906的粗体线所指示的，这些车道限制904和906可由系统100检测。车辆200可在车辆200与第一车道限制904相同的侧面上具有第一车辆侧面908，以及在车辆200与第二车道限制906相同的侧面上具有第二车辆侧面910。

[0157] 处理单元110可被配置为基于由图像捕获设备122-126采集的多个图像确定第一车道限制904和第二车道限制906，该多个图像可经由数据接口128由处理单元110接收。根据一些实施例，第一车道限制904和/或第二车道限制906能够以以上根据图8A-图8C所讨论的方式中的任意一种或全部来辨认、检测、估计或预测。

[0158] 如图所示，道路900可具有车道916。车道916可由车道限制界限。在示出的示例中，车道916在其左侧上由实线918标记，以及在其右侧上由实线120标记。在示出的示例中，车道916合并到车道902中作为在车辆200的行进的方向上继续的道路900。

[0159] 在图9A-图9C中，车道限制904在道路900上由实线903连续标记。另外，车道限制906被不一致地标记：在图的最底层部分，车道限制906在道路900上由实线905a标记；在该最底层部分以上的部分，随着车道916开始合并到车道902中，车道限制906在道路900上由虚线905b标记；在该部分以上的部分，随着车道916完成了合并到车道902中，车道限制906未被标记；最后，在图的最顶层部分，车道限制906再次由实线标记，这次在道路900上由实线905c标记。

[0160] 在图9A-图9C中，车道限制904、906和922通过系统100的检测由覆盖在车辆200的前面的相应的车道线的粗体线描绘。如图9A中的粗体线所描绘的，系统100在车辆200的前面的一定可见的距离内检测作为车辆200的左侧的实线车道限制的车道限制904。类似地，系统100在车辆200的前面的一定可见的距离内检测作为车辆200的右侧的实的车道限制的车道限制906。根据车道限制904和906的检测，系统100使车辆200在由车道限制904和906界定的车道902内行进。

[0161] 图9B描绘了在图9A中所描绘的时间点之后在车道902中沿着道路900继续前进的车辆200。随着车辆200到达车道916从右侧合并到车道902中的道路900的部分，其中车道限制906由虚线905b简要地标记并随后不由任何线标记，车道限制预测可由系统100执行。如图所示，系统100预测车道限制906的位置位于距相对的车道限制904的距离912处。在一些实施例中，该距离912可基于所检测到的历史信息以及由系统100储存的关于车道902的先前宽度和/或车道大致上的平均历史宽度来确定。如图所示，系统100基于在车道限制预测被激活之前车道限制906被系统100检测到的最后一次由虚线905b标记车道限制906的事

IPR2025-00943
Tesla EX1002 Page 1119

实，预测车道限制906继续是虚的车道限制。根据车道限制904的继续的检测和车道限制906的预测，系统100使车辆200在由车道限制904和906界定的车道902内继续行进。

[0162] 另外，图9B中示出的是检测车道限制922的系统100，该车道限制随着车道916合并到车道902中界限了车道916的右侧。系统100基于道路900上相应的实线920检测作为实的车道限制的车道限制922。在一些实施例中，系统100可确定被检测到的车道限制922不是车辆200行进的车道902的限制，而是合并到车道902中的车道的车道限制。在一些实施例中，该确定可根据车道限制922被检测到在接近的阈值速率以上接近车道902的确定来做出。例如，如果在车辆行进的车道的限制以外被检测到的车道限制使大于预定的阈值角度以上的角与车辆行进的车道的所检测的或先前所检测的车道限制平行，则可确定外部的车道限制对应于接近的合并车道。

[0163] 在一些实施例中，车道合并到另一车道中的确定可根据车辆中或其周围的道路中的其他方面、特性或对象的检测和分析来做出。例如，(例如，不同于提供系统的车辆的)第二车辆可被检测和跟踪，并且系统可确定另一车辆正在与道路上的车道成一角度地行进，即指示另一车辆行进的车道合并到另一车道中。在一些实施例中，系统可检测到另一车辆在接近的阈值速率以上从车道外的区域接近车道。在一些实施例中，系统可作出关于一个或多个其他车辆的这样的确定，并且可相应地确定在一个或多个其他车辆周围的车道合并。在一些实施例中，关于一个或多个其他车辆的信息可通过由与系统相关联的相机所捕获的多个图像来接收；在一些实施例中，可替代地或另外，关于一个或多个其他车辆的信息可通过任何合适形式的电子通信(诸如从一个或多个其他车辆中的一个或多个车辆传输的通信)来接收。

[0164] 在另外一些其他的实施例中，车道合并的确定可根据关于车道配置和/或布局的历史信息(诸如在车辆当前位置处或周围的车道的历史信息)来做出。

[0165] 图9C描绘了在图9B之后在车道902中沿着道路900继续前进的车辆200。当车辆200到达车道916合并到车道902中已经完成的道路900的部分时，由于车道限制906作为虚的车道限制又再次由系统100检测到(而不是被预测)，因此车道限制906又再次在道路900上由实线标记，这次由实线905c标记，并且车道限制预测可停止由系统100执行。在一些实施例中，系统100可通过确定车道限制不再接近在预定的接近的阈值速率以上的车辆(或车辆行进的车道)而确定其检测车辆200所在的车道的车道限制(而不是合并的车道的车道限制)。在一些实施例中，系统100可通过确定车道限制在如车道限制预测所预测的车道限制的所预测的位置的预定的阈值距离内确定其检测车辆所在的车道的车道限制(而不是合并的车道的车道限制)。

[0166] 如图9C中所描绘的，车辆200在执行车道限制预测时行进了距离914。在一些实施例中，车道限制预测仅可在预定的最大距离(和/或预定的最大时间)内被连续或基本上连续地执行。根据一些实施例，为车道限制预测所分配的最大距离/时间可被动态地设置。例如，为车道限制预测所分配的最大距离/时间可根据车辆的周围环境、车辆的速率、路的拓扑结构、亮度或环境、天气、道路上的其他车辆的密度和/或接近度、以及/或者任何其他路条件或特性来调整。在一些实施例中，当确定了接近的合并车道存在于车辆200的附近时，系统100可准许在比车道限制预测应另外准许执行更大的总的距离内执行车道限制预测。例如，虽然默认的最大距离可以是18米，但在检测到正在进行的车道合并的情况下，最大距

IPR2025-00943
Tesla EX1002 Page 1120

离可以是1000米。该差异可能是考虑到车道合并可能需要几百米来完成的事实的原因。在图8C所描绘的示例中，距离914小于用于在车道合并的情况下执行车道限制预测的预定的最大距离，使得车道限制预测能够随着车辆200横穿在其在图9B中的定位和其在图9C中的定位之间的距离被连续地执行。

[0167] 图10A和图10B图示了车辆200在道路1000上行进的实施例。图10A示出了如从车辆200的视角所见的三维视图，诸如当车辆200在道路1000上行进时由系统100接收的一个或多个图像中的一个。(图10A描绘了在车辆200的前面的道路1000上行驶的领先车辆，不要与车辆200本身混淆。)图10B示出了在道路1000的相同部分上的车辆200的架空图。

[0168] 在所描绘的示例中，可使用用于车道限制预测(例如，当存在护栏阴影条件时，预测车道限制的位置)和在由预测的车道限制界定的车道内操作车辆200的公开的系统和方法。如图所示，道路1000可具有车辆200行进的车道1002。车道1002可在第一侧面上由第一车道限制1004界定，以及在第一侧面相对的第二侧面上由车道限制1006界定。在一些实施例中，如示出车道限制1004和1006的粗体线所指示的，这些车道限制1004和1006可由系统100检测。车辆200可具有在车辆200与第一车道限制1004相同的侧面上的第一车辆侧面1008，以及车辆200在与第二车道限制1006相同的侧面上的第二车辆侧面1010。

[0169] 道路1000可包括车道线1003a和1003b，其分别在道路1000上物理地标记了车道1002的左侧和右侧。

[0170] 道路1000可包括由柱1014支撑的护栏1012。护栏1012和其柱1014二者都投射护栏阴影1016，该护栏阴影1016在所描绘的示例中落在车道线1003a上并落进车道1002中。

[0171] 如图10A和图10B中的粗体线所描绘的，系统100在车辆200的前面的一定可见的距离内检测作为车辆200的右侧的实线限制的车道限制1006。在道路1000的左侧，粗体线示出了系统1000如何检测一些边缘，其中的一个可归因于车道线1003a，以及其中的两个(阴影边缘1005a和1005b)可归因于护栏阴影1016。在一些实施例中，在车辆或车道的一侧上的多个边缘的检测可能使从系统100可用的图像中辨认车道限制更加困难，使得车道限制不能以高水平的确定性来辨认。因此，在一些实施例中，当系统100确定护栏1012或护栏柱1014存在时，即护栏阴影1016被投射到道路1000上，或者多个边缘在车道1002的相同侧面上被检测到，则系统100可确定护栏阴影条件存在，并且可相应地进入护栏阴影模式。在一些实施例中，护栏阴影模式可使对接近被检测的护栏、阴影或边缘的车道的侧面执行车道限制预测。通过使用车道限制预测，系统100可根据以上参照图8A-图8C和/或图9A-图9C所讨论的技术中的任意一种来预测车道限制1004的位置。在所描绘的示例中，系统100可正确地预测与由画出的车道线1003a界定的边缘而不是由护栏阴影1016界定的边缘一致的车道限制1002的位置。根据车道限制1004和1006的检测，系统100使车辆200在由车道限制1004和1006界定的车道1002内行进。

[0172] 在一些实施例中，护栏阴影预测可使车道限制预测以没有关于车道限制预测可执行多久的限制来进行。因此，例如，系统100可折向最大距离，而不是用于车道限制预测的最大距离被设置为18米(图8A-图8C的示例中)或1000米(图9A-图9C的示例中)，在功能上配置系统为使得只要继续检测到护栏阴影条件且只要没有检测到车道限制预测取消条件，则车道限制预测就可在无限的距离(和/或无限的时间)内继续。

[0173] 图11是存储器140或150的示例性功能框图，该存储器可储存用于执行与公开的实

IPR2025-00943
Tesla EX1002 Page 1121

施例一致的一个或多个操作的指令。如图11中所图示的，存储器140或150可储存用于执行本文中所描述的技术的一个或多个模块。例如，存储器140或150可储存图像获取模块1100、道路车道限制识别模块1102、确定模块1104、道路车道限制预测模块1106、车辆启动模块1108、和车道限制预测停止模块1110。如本领域中的普通技术人员将认识到的，以上列举的模块中的任意一个可与每个另外的模块或其他模块合并，可与另外一个模块或其他模块接合，以及/或者可被分成另外的模块和/或子模块。

[0174] 在一些实施例中，处理单元110可执行被储存在存储器140或150中的一个或多个模块上的指令，并且处理单元110的执行可使系统100执行本文中所讨论的方法。

[0175] 在一些实施例中，处理单元110可使系统100从相机(例如，图像捕获设备122、124和/或126)获得(例如，以图像获取模块1100)多个图像，其中，当相机安装在车辆(例如，车辆200)上时且在车辆处于运动状态下时，相机适于捕获在相机视场内的道路的图像，对多个图像应用图像处理以识别(例如，利用道路车道限制识别模块1102)第一道路车道限制，基于多个图像确定(例如，利用确定模块1104)一个或多个界定的车道预测条件被满足，根据一个或多个界定的车道预测条件被满足的确定，预测(例如，利用道路车道限制预测模块1106)第二道路车道限制的位置，以及根据第二道路车道限制的所预测的位置发布(例如，利用发布模块1108)通知或控制信号，以控制车辆的运动。

[0176] 在一些实施例中，处理单元110可使系统100根据第一道路车道限制的位置发布(例如，利用发布模块1108)通知或控制信号，以引导车辆的运动。

[0177] 在一些实施例中，处理单元110可使系统100发布(例如，利用发布模块1108)通知或控制信号，以沿着第一道路车道限制和第二道路车道限制之间的车道或区域引导车辆的运动。

[0178] 在一些实施例中，处理单元110可使系统100确定(例如，利用确定模块1104)一个或多个界定的车道预测取消条件是否满足，并根据一个或多个界定的车道预测取消条件被满足的确定，停止预测(例如，利用车道限制预测停止模块1110)第二道路车道限制的位置。

[0179] 图12说明了用于实施与公开的实施例一致的车道限制预测的过程1200。在一些实施例中，过程1200中的步骤可通过图8A-图8C、图9A-图9C和图10A-图10B中所图示的情况来进行说明。根据一些实施例，过程1200可由导航系统100的一个或多个组件(诸如至少一个处理单元110)实施。根据一些实施例，过程1200可由与车辆200相关联的系统100实施。

[0180] 在步骤1202，在一些实施例中，系统接收与系统相关联的移动的车辆附近的道路的多个图像。在一些实施例中，一个或多个图像使用至少一个图像捕获设备(诸如图像捕获设备122、124和/或126)被采集。多个图像可通过诸如数据接口128的数据接口由系统接收。在图8A-图8C中所图示的示例中，系统100可采集车辆200附近的道路800的多个图像。

[0181] 在步骤1204，在一些实施例中，系统识别第一道路车道限制。在一些实施例中，系统可使用图像分析技术，以通过辨认指示车道限制的图像中的对象、标记或其他特性(诸如车道标记、道路边缘、车道驾驶者和/或其他车辆)对用于检测车道限制所接收的一个或多个图像进行处理。在一些实施例中，系统可试图辨认车辆的一侧上的至少一个车道限制和车辆的另一侧上的另一车道限制(例如，在车辆的从前到后的行进的方向上，其中心线的每个侧面一个车道限制)。在一些实施例中，系统可试图辨认车道的一侧上的至少一个车道限制和车道的另一侧上的另一车道限制。在图8A-图8C中所图示的示例中，系统100可分析道

IPR2025-00943
Tesla EX1002 Page 1122

路800的所接收的图像，并可检测对应于实的车道线803a且与其相符的车道802的左侧的车道限制804。也就是说，在一些实施例中，车道线803a可由系统100在多个采集的图像中识别出来，并且系统100可确定车道线803a是用于车道限制804的视觉标记。

[0182] 在步骤1206，在一些实施例中，系统基于多个图像确定一个或多个界定的车道预测条件是否满足。在一些实施例中，车道预测条件可以是任意预定的条件，系统100通过该车道预测条件可确定是否预测车道限制的位置。在一些实施例中，仅一个车道预测条件需要满足，而在其他实施例中，车道预测条件的最小阈值数量(例如，两个或更多个)需要满足。在一些实施例中，车道预测条件的最小阈值数量可根据与道路车道布局、道路拓扑结构、阴影条件、能见度条件、沿着道路的护栏和/或其他障碍物的存在、道路上的其他车辆的存在等有关的历史信息来确定。

[0183] 在一些实施例中，一个或多个界定的车道预测条件包括随着车辆前进，基于多个图像，第二道路车道限制停止被识别。例如，在一些实施例中，系统可连续地捕获和查阅对应于车辆的当前位置和附近的多个图像；随着车辆沿着道路并沿着车道行进，基于多个图像的一个子集在一种情况下可被及时识别的车道限制可基于多个图像的后来的子集在随后的情况下及时停止被识别。在图8A-图8C的示例中，随着车辆200沿着道路800从其在图8A中的定位行进到其在图8B中的定位，系统100可分析道路800的所接收的图像。随着车辆200从其在图8A中的定位行进到其在图8B中的定位，该车辆200到达道路800的中间部分，在该中间部分有点的车道线803b未被画在道路800上。在对从道路800的该中间部分所接收的多个图像进行分析时，由于车道线803b的缺失，系统100可停止识别车道限制806。随着车辆200沿着道路800前进的该停止辨认车道限制在一些实施例中可构成车道预测条件。

[0184] 在一些实施例中，一个或多个界定的车道预测条件包括车道的第一侧面上的车道限制已经被辨认了至少阈值时间或阈值距离。可选地，与一个或多个界定的车道预测条件相关联的阈值时间或阈值距离可以是动态的，并且可根据与道路车道布局、道路拓扑结构、阴影条件、能见度条件、沿着道路的护栏或其他障碍物的存在、道路上的其他车辆的存在等有关的历史信息来调整。在一些实施例中，这样的车道预测条件可用于在系统第一次检测到在车道的单一侧面(或车辆的单一侧面)上的车道限制时，防止系统进行车道限制预测(或降低系统进行车道限制预测的可能性)。也就是说，在一些实施例中，如果车辆正在行进并出现第一次碰到单一车道线或者其他车道标记或指示物，则系统可基于单一车道标记/指示物不预测另一车道限制的位置。在一些实施例中，该车道预测条件可用于确保系统基于存在于道路上的实际车道而不是孤立的、杂散的或者错误或随机放置的车道标记/指示物作出有效且准确的车道限制预测(或增加其可能性)。在一些实施例中，预定的阈值距离可被固定在系统中，或者可基于车辆的当前速率、车辆行进的道路的大小或标示的速率、或与车辆或道路有关的任何其他可辨别的变量通过系统确定。

[0185] 在一些实施例中，一个或多个界定的车道预测条件包括护栏存在于车辆附近。在一些实施例中，当系统确定车辆在护栏附近时，系统可进行车道限制预测。在一些实施例中，该车道预测条件可用于防止当在现实中一个或多个边缘与护栏或与护栏阴影相关联时系统错误地确定在一个或多个图像中所检测到的一个或多个边缘对应于车道限制。在一些实施例中，护栏可将阴影投射到道路上，并且道路上的阴影的边缘可通过系统所使用的图像分析技术来检测。在一些实施例中，系统可能错误地确定护栏阴影是车道限制的标记，并

IPR2025-00943
Tesla EX1002 Page 1123

且可能根据护栏阴影而不是实际的车道限制使车辆行进。在一些实施例中，为了防止这样的错误，系统可识别护栏存在于车辆的附近，并因而可进行车道限制预测，以预测具有护栏的车道的侧面上的车道限制的位置，而不是试图识别车道的该侧面上的车道限制，并冒着风险错误地将护栏阴影识别为车道限制。

[0186] 在图10A和图10B的示例中，(图10B中所示的；不是图10A中所示的，因为车辆200在图10A中构成了视点)车辆200行进在道路1000上的车道1002中，该车道在车道1002的左侧具有护栏1012。在一些实施例中，系统100可分析一个或多个图像，以检测护栏1012存在于车道1002的左侧的车辆200的附近。在一些实施例中，护栏1012的存在可满足车道预测条件。

[0187] 在一些实施例中，确定护栏存在于车辆的附近包括在一个或多个图像中检测由支撑护栏的柱创建的图案。在一些实施例中，系统可被配置为在已知或怀疑指示护栏的存在的一个或多个图像中检测一个或多个图案。在一些实施例中，一个这样的图案可被称作“快速虚线(rapid dash)”图案，并且当系统在一个或多个图像中检测到由支撑护栏的柱产生的快速连续的虚线或快速连续的线时，这样的图案可能存在。在一些实施例中，系统可被配置为检测已知为支撑护栏的柱的特性的任何其他合适的图案或视觉记号。在一些实施例中，系统可检测护栏的柱的图案，而在一些实施例中，可替代地或另外地，系统可检测由护栏的柱投射的阴影的图案。在图10A和图10B的示例中，护栏1012由(图10A中所示的；从图10B的架空角(overhead angle)不可见的)柱1014支撑。在图10A和图10B中，柱1014投射阴影，该阴影落在道路1000上并落在车道1002中。系统100可检测快速虚线图案或在多个图像中由柱1014或其阴影创建的任何其他合适的图案的存在，并且该检测可以是系统100确定护栏存在于车辆200的附近的一种形式。

[0188] 在一些实施例中，系统100可利用深度信息来辨认是护栏或护栏的部分的特性的形状。应认识到的是，画出的线将基本上没有体积，并且其高度相对于道路表面几乎为零，因此，深度信息可用于在车道标记和护栏之间进行区分。

[0189] 在一些实施例中，一个或多个界定的车道预测条件包括护栏将阴影投射到道路上。在一些实施例中，代替通过系统对护栏本身的检测或除此之外，车道预测条件还可基于通过系统对护栏的阴影的检测。在一些实施例中，如果护栏阴影被检测为在车辆的附近的任何地方，则车道预测条件可被满足；在一些实施例中，如果阴影被检测到为来自车辆行进于其中或其上的车道或道路的方向上的护栏的投射，则车道预测条件可被满足；在一些实施例中，如果阴影被检测到为从护栏投射到车辆行进于其中的道路或车道上，则车道预测条件可被满足。

[0190] 在一些实施例中，确定护栏是将阴影投射到道路上可涉及确定预测的太阳的位置是否与可能是护栏阴影的所检测到的边缘的位置一致。例如，在一些实施例中，系统可确定太阳相对于可用于系统的多个图像的定向和太阳的定位。这些关于太阳的位置和车辆的定向的确定可根据各种信息碎片(例如，包括多个图像中的一个或多个、由GPS确定的车辆的位置和/或定向、由罗盘确定的车辆的定向和/或时间)来完成。系统可基于所确定的太阳的位置和车辆的定向，确定阴影预计在某个时间某个方向上从护栏投射。

[0191] 在图10A和图10B的示例中，护栏1012是将阴影1016投射到道路1000上并投射到车道1002中；落在道路1000上并落在车道1002中的阴影1016可由系统100检测，并且所述检测

IPR2025-00943
Tesla EX1002 Page 1124

可满足车道预测条件。

[0192] 在一些实施例中，一个或多个界定的车道预测条件包括多个基本上平行的边缘在多个图像中被检测到在道路上的车道的一侧上的道路上。在一些实施例中，由系统应用到所接收的多个图像的图像分析技术可在所接收的多个图像中的一个或多个中检测边缘；在一些实施例中，边缘的检测可用于识别车辆周围的环境中的对象和标记。在一些实施例中，多个图像中的边缘的检测可用于识别车道线或标记或边缘或道路，并且从而用于确定车道限制的位置。在一些实施例中，包括护栏阴影被投射到车辆行进的道路或车道上的那些实施例，系统可在多个图像中检测作为边缘的护栏阴影的边缘。在一些其他的实施例中，其他对象或标记可使系统检测一个车道的相同侧面上的多个边缘。在这些实施例中，哪个边缘对应于正确的车道限制的确定可能是困难的、缓慢的或者不可靠的。出于这个原因，在一些实施例中，可能有利的是，不基于多个图像来试图确定几个所检测到的边缘中的哪个边缘对应于车道限制，并如以上所讨论的，当多个边缘在多个图像中被检测到时，使用车道限制预测代替简单地预测车道限制的位置。在图10A和图10B的示例中，系统100在多个图像中检测车道1002的左侧由粗体线标记的1004、1005a和1005b指示的三个不同的边缘。在图10A和图10B的示例中，如以上所解释的，系统100可确定多个边缘在车道1002的左侧被检测到，并且该检测可构成车道预测条件。

[0193] 在一些实施例中，如果在框1206处确定一个或多个界定的车道预测条件未被满足，则方法1200可返回到步骤1202，并且可随着时间的推移继续采集道路的额外的图像。

[0194] 在一些实施例中，如果在框1206处确定一个或多个界定的车道预测条件满足，则方法1200可前进到步骤1208。在步骤1208，在一些实施例中，根据一个或多个界定的车道预测条件被满足的确定，系统预测第二道路车道限制的位置。在一些实施例中，系统可预测车道的与所检测到的车道限制的相对侧面上的车道限制(车辆作为相对侧面的方向的参考)。因此，在一些实施例中，如果一个车道限制可从所接收的图像中的车道线或其他视觉记号识别出，并且确定诸如以上所讨论的条件中的任意一个的车道预测条件满足，则系统可预测车道的与第一车道限制相对侧面上的第二车道限制的位置。

[0195] 在一些实施例中，所预测的车道限制可被预测为沿着道路在预定的距离内延伸。在一些实施例中，所预测的车道限制可被预测为平行于与其相对的所识别的车道限制。在一些实施例中，所预测的车道限制可被预测为继续沿着与先前由系统检测的相同的车道限制相同的角度、路径或位置。在一些实施例中，所预测的车道限制可基于另一车道限制的位置、基于道路的其他可辨别的特征的位置、或基于与系统和/或周围的车辆相关联的车辆的位置来预测。在一些实施例中，所预测的车道限制可被预测为位于距被识别的车道限制或距基于道路的多个图像所识别或预测的任何其他对象的预定的距离处。在一些实施例中，所预测的车道限制可被预测为是匹配在多个图像中所检测的其他车道限制类型或根据其确定的车道限制类型，或者被预测为是根据基于从先前时间的道路的先前的图像对所预测的车道限制所确定的车道限制类型确定的车道限制类型(例如，如果系统在车辆的右侧检测到实的单一的白线，则停止检测任何车道线，系统可预测实的单一白的车道限制在车辆的右侧继续；可替代地或另外地，系统可预测实的单一白线位于与实的单一黄线相对的车道的侧面上)。

[0196] 在图8A-图8C的示例中，系统100基于多个图像中的车道线803a的检测识别车道限

IPR2025-00943
Tesla EX1002 Page 1125

制804。当车辆200到达其在图8B中的定位且没有在车道802的右侧检测到任何车道线时，则车道预测条件可被满足。因此，在一些实施例中，系统100可预测车道802的右侧的车道限制806，如图8B中所示。在一些实施例中，车道限制806可基于以是单一虚线车道限制的车道限制806的车道线803b为基础的图8A中的先前的识别而被预测为单一虚线车道限制。在一些实施例中，车道限制806的定位、角度、类型和/或任何其他特性可基于通过系统100可访问的关于车辆200、道路800或任何其他周围环境的因素的任何储存的或以其他方式确定的信息来预测。

[0197] 在一些实施例中，预测车道限制的位置包括基于关于先前确定的车道的宽度预测车道的宽度。在一些实施例中，车道的宽度可基于关于车道宽度的历史信息或关于最新(例如，在执行车道限制预测之前的几秒钟)的特定车道的宽度的缓存信息来预测。在一些实施例中，当系统检测车道的两个限制并且从而可确定该车道的宽度并且系统随后停止检测或识别车道限制中的一个时，则系统可预测车道继续是相同的宽度。系统可相应地预测所预测的车道限制的位置，使得车道被预测为保持宽度恒定或基本上恒定。在图8A-图8C的示例中，当系统100预测图8B中的车道限制806的位置时，限制806被预测为离车道限制804距离812。当车辆200处于图8A中所示的定位时，该距离812可以是车道限制806和车道限制804之间所检测到的相同的距离812。在一些实施例中，系统100可记录或缓存关于车道宽度的历史信息，使得该系统可快速调回车道宽度信息，以便对所预测的车道限制的位置进行预测。

[0198] 在步骤1210，在一些实施例中，系统使移动的车辆能够避免第一道路车道限制和第二道路车道限制。因此，在一些实施例中，系统可使车辆在所检测/识别的车道限制在一侧上界定的和所预测的车道限制的相对侧面上所界定的车道中行进。在一些实施例中，系统可引起车辆的任何合适的转向、加速、减速或其他控制，以使车辆在由所预测和检测的车道限制界定的车道内行进。在图8A-图8C的示例中，当系统100检测到图8B中的车道限制806时，系统100随后使车辆200沿着所检测到的车道限制804和所预测的车道限制806之间的车道802行进，以从其在图8B中的定位行进到其在图8C中的定位。

[0199] 在步骤1212，在一些实施例中，系统像确定一个或多个界定的车道预测取消条件是否被满足。在一些实施例中，车道预测取消条件可以是任意预定的条件，系统通过该车道预测取消条件可确定是否停止预测车道限制的位置。在一些实施例中，一旦系统开始预测车道限制的位置，车道限制的位置就可被继续预测，直至一个或多个预定的车道预测取消条件被满足为止。在一些实施例中，仅一个车道预测取消条件需要被满足，以便停止车道预测。在一些实施例中，预定的数量或车道预测条件大于必须满足的一个条件，以便停止车道预测。可选地，置信分数可根据预定的标准分配给车道预测条件。可选地，一些车道预测条件可具有比其他更高的权重。仍进一步通过示例的方式，可计算总体分数，以确定车道预测取消条件是否满足，并且这样的计算可考虑被分配给车道预测条件的权重和置信分数。

[0200] 在一些实施例中，一个或多个界定的车道预测取消条件包括在停止识别之后，基于第一或多个图像识别第二车道限制。例如，在一些实施例中，随着车辆沿着道路行进，车道限制可停止被识别(如以上所讨论的)，诸如当道路未被标记或车道线在多个图像中不可检测时；如以上所讨论的，可相应地执行车道预测。在一些实施例中，此后，例如当车辆继续沿着道路行进时，随着系统继续采集新的图像，系统可在多个图像中再次识别到车道限制。例如，如果车辆到达车道线不再遗漏的道路的部分，或者如果车道线以其他方式变成可由

IPR2025-00943
Tesla EX1002 Page 1126

系统再次检测到，则系统可基于道路的更新的图像来识别车道限制。该第二车道限制的识别-被预测的车道限制-在一些实施例中可构成车道预测取消条件，该车道预测取消条件可使系统停止预测第二车道限制的位置并转而返回依赖于基于多个图像所检测/识别的车道限制的位置。

[0201] 在图8A-图8C的示例中，虽然系统100可在车辆200位于图8B中所示的位置时预测车道限制806的位置，但随着车辆200前进到图8C中所示的定位，图的上半部分上的虚的车道线803在由系统100采集的图像中可再次是可见的。基于这些新的图像，系统100可基于车道线803b再次检测/识别车道限制806的存在，并且该识别可构成车道预测取消条件，该车道预测取消条件触发车道限制806的位置预测的停止。

[0202] 在一些实施例中，一个或多个界定的车道预测取消条件包括驾驶者在车道之间移动车辆。在一些实施例中，系统可基于多个图像确定驾驶者将车辆转向或合并到另一车道。在一些实施例中，该检测可满足车道预测取消条件，并且可触发车道限制位置预测的取消。

[0203] 在一些实施例中，一个或多个界定的车道预测取消条件包括车辆横穿超过预定的锐度阈值的曲线。在一些实施例中，系统可访问关于曲线的预定或动态确定(例如，根据车辆速率、标示的限速、道路状况、天气状况、环境亮度、其他车辆的附近等确定)的锐度阈值，并且可确定道路或车道中的曲线是否超过锐度阈值(例如，比最大允许的锐度更锐利)。如果确定了车辆横穿或接近的曲线超过最大允许的锐度，则系统可确定车道预测取消条件满足，并且可相应地停止车道限制预测。

[0204] 在一些实施例中，确定一个或多个界定的车道预测取消条件是否满足包括选择最大车道预测距离设置，并确定车辆在预测第二道路车道限制的位置时所行进的距离是否大于与最大车道预测距离设置相关联的最大车道预测距离。在一些实施例中，最大车道预测距离可由系统预定、设置或动态地确定。在一些实施例中，系统可监视车辆在系统预测车道限制的位置时所行进的距离，并且如果该行进的距离超过最大允许的车道预测距离，则可停止预测车道限制的位置。在一些实施例中，最大允许的车道预测距离的实施可通过确保系统在过长的时间段或过长的距离内不预测车道限制的位置来改进安全性和功能，在该距离或时间内，车道或道路的一个或多个特性可能改变或变化。在一些实施例中，由于可预期车道和道路的特性随着空间和时间而改变或变化，因此系统可通过确保至少一旦超过最大车道预测距离(或者，在一些实施例中，是最大车道预测时间)车道限制参考点或指示物(诸如车道线或其他车道标记)就基于由系统采集的多个图像来进行检测而改进安全性。

[0205] 在一些实施例中，最大车道预测距离在不同的情况下可由系统改变。在一些实施例中，最大车道预测距离可与一个或多个最大车道预测距离设置相关联，使得系统可选择一个最大车道预测距离设置并实施相关联的最大车道预测距离。在一些实施例中，不同的最大车道预测距离设置可在由系统检测和确定的不同的驾驶场景中实施。

[0206] 在一些实施例中，标准或默认的最大车道预测距离设置可与标准或默认的最大车道预测距离相关联。在一些实施例中，标准或默认的最大车道预测距离可以是18米。

[0207] 在图8A-图8C的示例中，车辆200从并车道限制预测的其在图8B中的定位行进距离814到车道预测停止的、其在图8C中的定位。该距离814可小于18米的标准或默认的最大车道预测距离，使得车道限制预测可在车辆200在图8B和图8C中的定位之间被连续地执行。

[0208] 在一些实施例中，选择最大车道预测距离设置包括，根据第一车道与第二车道在

IPR2025-00943
Tesla EX1002 Page 1127

车辆的附近合并的确定，选择最大车道预测距离设置，使得该最大车道预测距离大于默认的最大车道预测距离。因此，在一些实施例中，系统可实施延伸的或增加的最大车道预测距离。例如，在一些实施例中，系统可基于多个图像确定车道与车辆行进的车道合并，或者车辆行进的车道与另一车道合并。系统可通过检测临近车辆的车道的合并的车道的远侧上的车道限制(例如，通过识别车道线)以及通过确定临近车道的相对侧面上的车道限制在预定或动态确定的接近的阈值速率以上接近车辆行驶的车道，来确定车道合并发生在车辆的附近。也就是说，在一些实施例中，随着车辆沿着道路前进，系统可识别接近车辆行驶的车道的临近车道的远侧，并且可相应地确定临近车道与车辆行驶的车道合并。

[0209] 在一些实施例中，当两个车道彼此合并时，预计道路可能将以长距离的无标记的车道限制为特征，因此，对于被配置为在比默认的最大车道预测距离或时间更大的距离和/或更长的时间内预测车道限制的位置的系统来说可能是有利的。在一些实施例中，与车道合并状况相关联的增加的最大车道预测距离可以是1000米。

[0210] 在图9A-图9C的示例中，车辆200从车道限制预测开始的、其在图9B中的定位到车道预测停止的、其在图9C中的定位行进了距离914。距离914可大于标准或默认的最大车道预测距离，使得如果选择了默认的最大车道预测距离设置，则在车辆200到达其在图9C中的定位之前，车道限制预测可已经被停止。然而，如图8B中的粗体线所指示的，随着车辆200前进，系统100识别接近车道902的车道916的右侧的车道限制922，并且可相应地确定车道正在合并，并应选择与车道合并相关联的最大车道预测距离设置。所选择的设置可与大于18米的默认的最大车道预测距离的1000米的最大车道预测距离相关联。因此，在一些实施例中，距离914可小于1000米的增加的最大车道预测距离，使得车道限制预测可在图9B和图9C中车辆200的定位之间被连续地执行。

[0211] 在一些实施例中，选择最大车道预测距离设置包括，根据界定的护栏条件满足的确定，选择没有最大车道预测距离的最大车道预测距离设置。也就是说，在一些实施例中，当系统100检测到护栏条件满足时，系统100可无限地预测车道限制的位置，而无需考虑任何最大车道预测距离。通过这种方式，在一些实施例中，与护栏条件相关联的最大车道预测距离可被认为是无限的最大距离或不存在的最大距离。

[0212] 在一些实施例中，护栏条件可在系统检测车辆附近的一个或多个护栏、护栏柱和/或护栏阴影的存在的任何情况下得到满足。例如，护栏条件可在以上所讨论的任意一种或全部情况下得到满足，其中，系统可确定护栏和/或护栏阴影存在于车辆的附近和/或车辆行进的道路或车道上。如以上所讨论的，在一些情况下，当护栏和/或护栏阴影的存在、位置和/或定向使系统进行车道限制预测时，护栏条件可被满足。

[0213] 在一些实施例中，对于不将最大车道预测距离(或无限的最大车道预测距离)应用到护栏和/或护栏阴影存在的情况中可能是有利的，因为预计护栏和其相关联的阴影沿着道路可连续持续非常长的距离(通常大于默认的最大车道预测距离或甚至大于与车道合并相关联的增加的车道预测距离)。在图10A和图10B的示例中，当由于护栏阴影1016的检测系统100开始预测车道限制1004的位置时，系统100可选择与护栏条件相关联的最大车道预测距离设置；与护栏条件相关联的最大车道预测距离设置的选择可使系统执行车道限制1004的车道预测，而无需以任何最大车道预测距离为前提。

[0214] 在一些实施例中，如果在框1212处确定一个或多个界定的车道预测取消条件未满

IPR2025-00943
Tesla EX1002 Page 1128

足，则方法1200可返回到步骤1202，并且可随着时间的推移继续采集道路的额外的图像。

[0215] 在一些实施例中，如果在框1212处确定一个或多个界定的车道预测取消条件满足，则方法1200可前进到步骤1214。在步骤1212，在一些实施例中，根据一个或多个车道预测取消条件满足的确定，系统停止预测第二车道限制的位置。在图8A-图8C的示例中，当车辆200到达其在图8C中所描绘的定位时，车道预测取消条件可基于车道线803b通过车道限制806的识别得到满足，并且因此系统100可停止预测车道限制806的位置。相反，如系统100先前在图8A中所做的，其可基于所检测到的车道线803b的位置重新确定车道限制806的定位。

[0216] 施工区的检测与响应

[0217] 系统100可给车辆200提供施工区检测，使得系统识别道路上的施工区的存在并对其采取动作。在一些实施例中，系统100可在其接收的一个或多个图像中识别各种元素，并可根据这些一个或多个元件的识别确定车辆200接近施工区或在其中。在一些实施例中，可根据施工区的检测采取任何自动或驾驶者辅助的动作，包括但不限于改变车辆200的速率、转向车辆200、启用或解除车辆200的一个或多个灯等。在一些实施例中，根据施工区的检测，信号可由系统100输出，该系统100指示施工区已经被检测到和/或施工区模式被激活。在一些实施例中，根据施工区的检测(和/或根据复杂的施工区的检测)，由系统100进行的对车辆200的自动控制或驾驶者辅助的控制可被禁用，使得车辆200的控制可返回给驾驶者。

[0218] 图13图示了车辆200行进在道路1300上并接近施工区的实施例。道路1300具有两个车道1302和1304。车道1302由车道限制1306和1308界定，而车道1304由车道限制1310和1312界定。如指示车道限制1306、1308、1310和1312的粗体线所示，这些车道限制由系统100检测。车道限制1306和1308分别通过画出的车道线1320和1322被标记在道路1300上，而车道限制1310和1312分别通过画出的车道线1322和1324被标记在道路1300上。

[0219] 道路1300包括施工区或施工区域的几个指示物，包括画出的施工区车道线1326。在一些实施例中，施工区车道线可用与传统车道线不同的颜色和/或式样画出。例如，在德国，黄色的车道线被画在路上来指示施工区，而传统车道线用白色画出。如粗体线所指示的，系统100基于画出的施工区车道线1326检测施工区车道限制1314。根据施工区车道限制1314的检测，系统100可使车辆200在由施工区车道限制1314而不是传统车道限制界定的车道内行驶。在一些实施例中，当系统100检测到车辆200在施工区中时，施工区车道线可优先于传统车道线。

[0220] 道路1300可包括传统施工区指示物，包括圆锥体1316，施工区指示物中的一些位于道路1300的表面上和车道1304内。响应于检测圆锥体1316，系统100可使车辆200改变其路径，以便避免圆锥体1316，如有必要改变行进的车道中的定位或改变车道。响应于检测圆锥体1316，系统100可确定车辆200在施工区中或在其附近；激活施工区模式；输出施工区信号；禁用车辆200的自动控制或驾驶者辅助控制；以及/或者通过改变其速率、转向或采取任何其他适当的动作(诸如避免圆锥体1316)控制车辆200。

[0221] 道路1300可包括额外的施工区指示物，包括施工路标志1318a和1318b，其可指示施工区的存在和/或为驾驶者指示在施工区中要遵循的指令(例如，减速、停止、合并、转弯等)。系统100可通过一种方式检测施工路标志1318a和1318b，在该方式中系统可检测任何

IPR2025-00943
Tesla EX1002 Page 1129

或所有其他的交通标志，并且该系统可响应于施工路标志1318a和1318b的检测采取任何适当的动作，包括确定车辆200在施工区中或在其附近；激活施工区模式；输出施工区信号；禁用车辆200的自动控制或驾驶者辅助控制；以及/或者通过改变其速率、转向或采取任何其他适当的动作(诸如根据施工路标志上的指令)控制车辆200。

[0222] 图14是存储器140或150的示例性框图，该存储器可储存用于执行与公开的实施例一致的一个或多个操作的指令。如图14中所图示的，存储器140或150可储存用于执行本文中所描述的施工区检测的一个或多个模块。例如，存储器140或150可储存图像获取模块1400、施工区指示物识别模块1402、确定模块1404、信号输出模块1406、车辆移动启用模块1408、和车辆控制禁用模块1410。如本领域中的普通技术人员将认识到的，以上列举的模块中的任意一个可与每个另外的模块或其他模块合并，与另外一个模块或其他模块接合，以及/或者被分成另外的模块和/或子模块。

[0223] 在一些实施例中，处理单元110可执行被储存在存储器140或150中的一个或多个模块上的指令，并且处理单元110的执行可使系统100执行本文中所讨论的方法。

[0224] 在一些实施例中，处理单元110可使系统100从相机(例如，图像捕获设备122、124和/或126)获得(例如，以图像获取模块1400)多个图像，其中，当相机安装在车辆(例如，车辆200)上且当车辆处于运动状态下时，该相机适于捕获其视场内的道路的图像，对多个图像应用图像处理以识别(例如，利用施工区指示物识别模块1402)施工区的一个或多个指示物，以及，根据识别施工区的一个或多个指示物，输出(例如，利用信号输出模块1406)输出指示车辆接近施工区的信号。

[0225] 在一些实施例中，处理单元110可使系统100基于一个或多个指示物确定(例如，利用确定模块1404)车辆接近施工区，并根据车辆接近施工区的确定，输出(例如，利用信号输出模块1406)指示车辆接近施工区的信号。

[0226] 在一些实施例中，处理单元110可使系统100根据车辆接近施工区的确定使移动的车辆(例如，利用车辆移动启用模块1408)根据施工区指示物移动。

[0227] 在一些实施例中，处理单元110可使系统100基于一个或多个指示物确定(例如，利用确定模块1404)施工区是复杂的施工区，并根据施工区是复杂的施工区的确定，禁用(例如，利用车辆控制禁用模块)自动车辆控制功能。

[0228] 图15说明了用于实施与公开的实施例一致的车道限制预测的过程1500。在一些实施例中，过程1500的步骤可通过图13中所图示的情况来进行说明。根据一些实施例，过程1500可由导航系统100的一个或多个组件(诸如至少一个处理单元110)实施。根据一些实施例，过程1500可由与车辆200相关联的系统100实施。

[0229] 在步骤1502，在一些实施例中，系统接收与系统相关联的移动的车辆附近的道路的多个图像。在一些实施例中，一个或多个图像使用至少一个图像捕获设备(诸如图像捕获设备122、124和/或126)被采集。多个图像可通过诸如数据接口128的数据接口由系统接收。在图13中所图示的示例中，系统100可采集车辆200附近的道路1300的多个图像。

[0230] 在步骤1504，在一些实施例中，系统在多个图像中识别施工区的一个或多个指示物。在一些实施例中，系统可使用图像分析技术来对所接收的一个或多个图像进行处理，以检测标志、对象、其他车辆、标记的车道线和周围环境的任何其他对象或特性。在一些实施例中，某些标志、标记、车道线、对象和/或任何这样检测的方面的组合可由系统预定为施工

IPR2025-00943
Tesla EX1002 Page 1130

区的指示物。在一些实施例中，圆锥体、圆筒体、闪光、杆、障碍物和/或存在于道路附近的标记可由系统预定为施工区的指示物。在一些实施例中，预定的交通标志，诸如那些承载的文字、符号或对于指示施工区已知的图像，可由系统预定为施工区的指示物。在一些实施例中，车道线可由系统预定为施工区的指示物。例如，在德国，施工区可用画在道路上的专用施工区车道线来标记，施工区车道线可能是与传统车道线不同的颜色或不同的外观。系统可检测车道线的颜色、外观、定向和/或位置，其向系统指示该线是施工区车道线，并且系统可识别这样的施工区车道线为施工区的指示物。

[0231] 在图13的示例中，系统100可识别圆锥体1316、施工路标志1318a和1318b、和/或画出的施工区车道线1326。在一些实施例中，这些元素中的任意一个或全部可由系统100预定或由系统100动态确定为施工区指示物。

[0232] 在步骤1506，在一些实施例中，系统基于一个或多个指示物确定车辆是否接近施工区。在一些实施例中，系统可基于施工区指示物的(车辆、道路或车道的)附近确定车辆接近施工区。在一些实施例中，系统可基于相对于道路或车道的定向或位置假设车辆接近施工区。在一些实施例中，系统可根据一个或多个路标志的内容(例如，图像、符号和/或文字)确定车辆接近施工区。在一些实施例中，最小阈值数量的施工区指示物必须在彼此的最大阈值时间内被检测，以便系统确定车辆接近施工区。在图13的示例中，系统100可基于圆锥体1316、施工路标志1318a和1318b和/或画出的施工区车道线1326的同时检测确定车辆200接近施工区。

[0233] 在一些实施例中，如果在框1506处确定车辆不接近施工区，则方法1500可返回到步骤1502，并且可随着时间的推移继续采集道路的额外的图像。

[0234] 在一些实施例中，如果在框1506处确定车辆接近施工区，则方法1500可前进到步骤1508。在步骤1508，在一些实施例中，根据车辆接近施工区的确定，系统输出指示车辆接近施工区的信号。在一些实施例中，由系统100输出的信号可以是电信号或电消息，其向系统的任何组件或向任何相关联的系统通知施工区被检测到。例如，信号可被发送到与系统相关联的显示器系统，或者信号可被发送到与系统在通信上耦合的任何其他本地系统或远程系统。在一些实施例中，信号可由车辆的驾驶者察觉。例如，信号可以是可见的信号(诸如计算机或娱乐系统显示器上的发光指示物或显示元素)、由扬声器系统发出的可听信号和/或诸如振动的触觉信号。

[0235] 在步骤1510，在一些实施例中，根据车辆接近施工区的确定，系统使移动的车辆能够根据施工区指示物移动。在一些实施例中，响应于车辆接近施工区的确定，系统可使车辆根据所检测到的施工区指示物加速、减速、转向或以其他方式受控和移动。在一些实施例中，车辆可自动减速到降低的速率，以增加安全性。在一些实施例中，可使车辆转向或改变车道，以便跟随施工区车道线，避免圆锥体/圆筒体/杆/障碍物或其他施工区指示物，或者跟随在指示施工区或与其相关联的路标志上所指示的指令。在图13的示例中，系统100可修改车辆200的行进的进程，使得车辆200根据施工区车道线1326跟随所检测到的车道限制1314(而不是根据车道线1324所检测到的车道限制1312)，避开圆锥体1316并遵守交通标志1318(其包括指示车辆200向左合并的符号)。在一些实施例中，可使车辆200相应地从车道1304向左合并到车道1302中。

[0236] 在步骤1512，在一些实施例中，系统基于一个或多个指示物确定施工区是否是复

IPR2025-00943
Tesla EX1002 Page 1131

杂的施工区。在一些实施例中，复杂的施工区可以是具有一个或多个预定的或动态确定的特性的施工区，该预定的或动态确定的特性通过系统与施工区相关联，在该施工区中车辆的自主控制和/或驾驶者辅助控制可能是不安全的，其中，车辆的自主控制和/或驾驶者辅助控制可能是不可靠的或不可信赖的，或者其中，车辆的控制应返回给车辆的驾驶者。在一些实施例中，复杂的施工区可以是车辆行驶于其上或其中的车道中或道路上存在阻碍的施工区。在一些实施例中，复杂的施工区可以是预定或动态确定的阈值数量的圆锥体/圆筒体/杆/障碍物或其他施工区指示物在彼此的阈值距离或阈值时间内被检测到的施工区。也就是说，在一些实施例中，当施工区被密集填充了施工区指示物时，其可被确定为复杂的施工区。在图13的示例中，在一些实施例中，道路1300可基于道路1300上的圆锥体1316的存在、车道1304中的圆锥体1316、所检测到的圆锥体1316的最小阈值数量、和所检测到的总的施工区指示物的最小阈值数量，由系统100确定为包含复杂的施工区。

[0237] 在一些实施例中，如果在框1512处确定施工区不是复杂的施工区，则方法1500可返回到步骤1502，并可随着时间的推移继续采集道路的额外的图像。

[0238] 在一些实施例中，如果在框1512处确定施工区是复杂的施工区，则方法1500可前进到步骤1514。在步骤1514，在一些实施例中，根据施工区是复杂的施工区的确定，系统可禁用自动车辆控制功能。在一些实施例中，驾驶者辅助功能可被禁用。在一些实施例中，完全自动驾驶功能可被禁用。在一些实施例中，车辆200的全部控制(例如，加速、减速、转向等)可由系统100返回给车辆200的驾驶者。

[0239] 出于说明的目的呈现了以上描述。该描述并不详尽，但不限于公开的精确形式或实施例。从公开的实施例的规范和实践的考虑，修改和调整对于本领域中的那些技术人员来说将是明显的。另外，虽然公开的实施例的方面被表述为储存在存储器中，但本领域中的技术人员将认识到，这些方面也可储存在诸如二级储存设备的其他类型的计算机可读媒介中，例如，硬盘或CD ROM，或者其他形式的RAM或ROM、USB媒介、DVD、Blu-ray、或其他光驱动媒介。

[0240] 基于书面的描述和公开的方法的计算机编程在有经验的开发人员的技能的范围内。各种程序或程序模块可使用与本领域中的技术人员已知的任何技术来创建，或者可结合现有的软件来设计。例如，程序段或程序模块能够采用或通过.Net Framework、.Net Compact Framework(和有关的语言，诸如Visual Basic、C等)、Java、C++、面向对象的C、HTML、HTML/AJAX的组合、XML、或包括Java小程序的HTML的方法来设计。此外，应理解的是，本文中所描写的程序模块的功能不限于本文中所命名的特定类型的程序模块，而是可在一个或多个模块或其他软件结构之间以任何合适的方式来编程。

[0241] 另外，虽然已经在本文中对说明性的实施例进行了描述，但实施例中的任意一个实施例或全部实施例的范围具有如本领域中的那些技术人员基于本公开将所认识到的等效的元素、修改、省略、组合(例如，跨各个实施例的方面的组合)、调整和/或改变。权利要求中的限制基于权利要求中所采用的语言被广泛解释，并且不限于在本说明书或在申请的贯彻期间所描述的示例。示例被解释为非排他性。此外，公开的方法的步骤能够以任何方式修改，包括通过重新排序步骤以及/或者插入或删除步骤。因此，意图是说明书和示例被认为仅用于说明，真实范围和精神由下面的权利要求及其等效形式的全部范围指示。

IPR2025-00943
Tesla EX1002 Page 1132

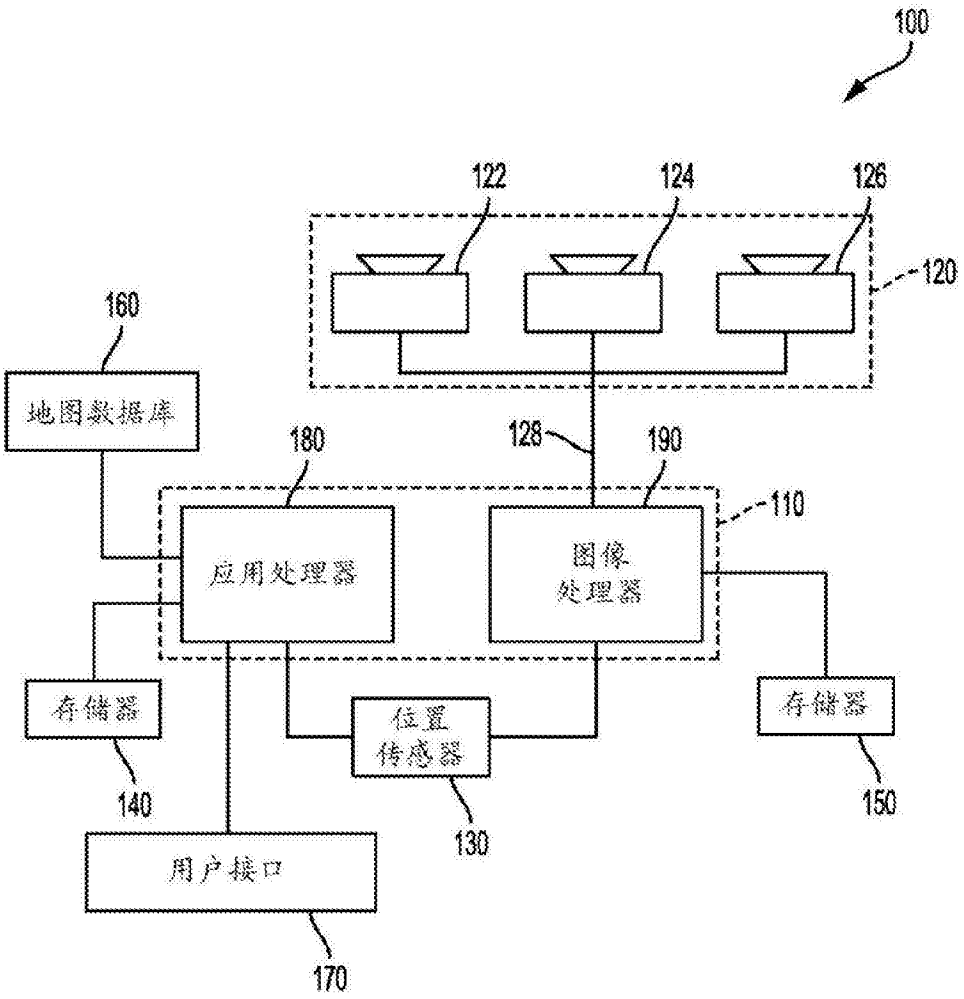


图1

IPR2025-00943
Tesla EX1002 Page 1133

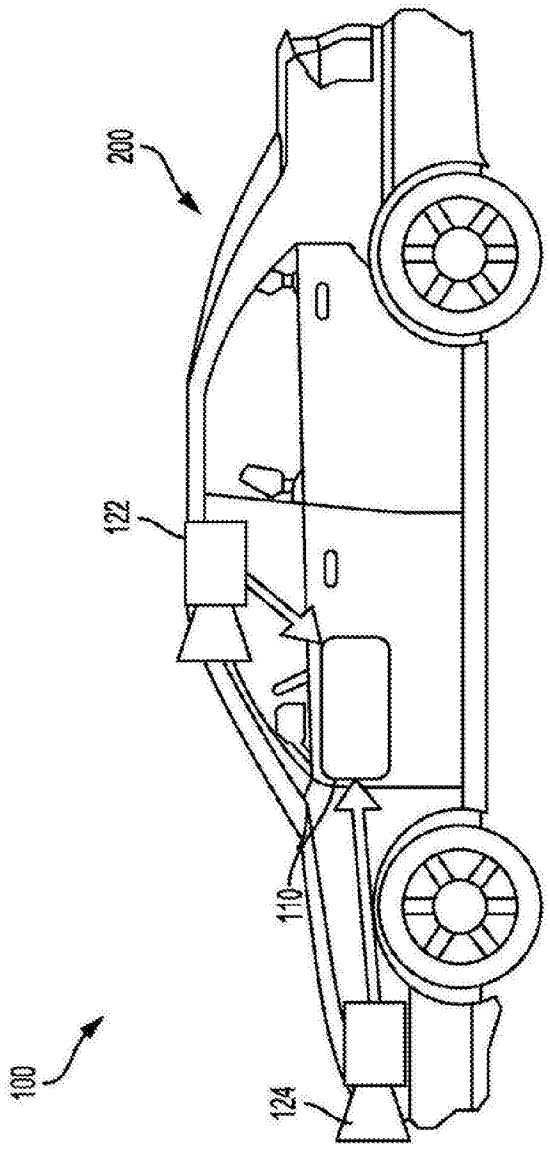


图2A

IPR2025-00943
Tesla EX1002 Page 1134

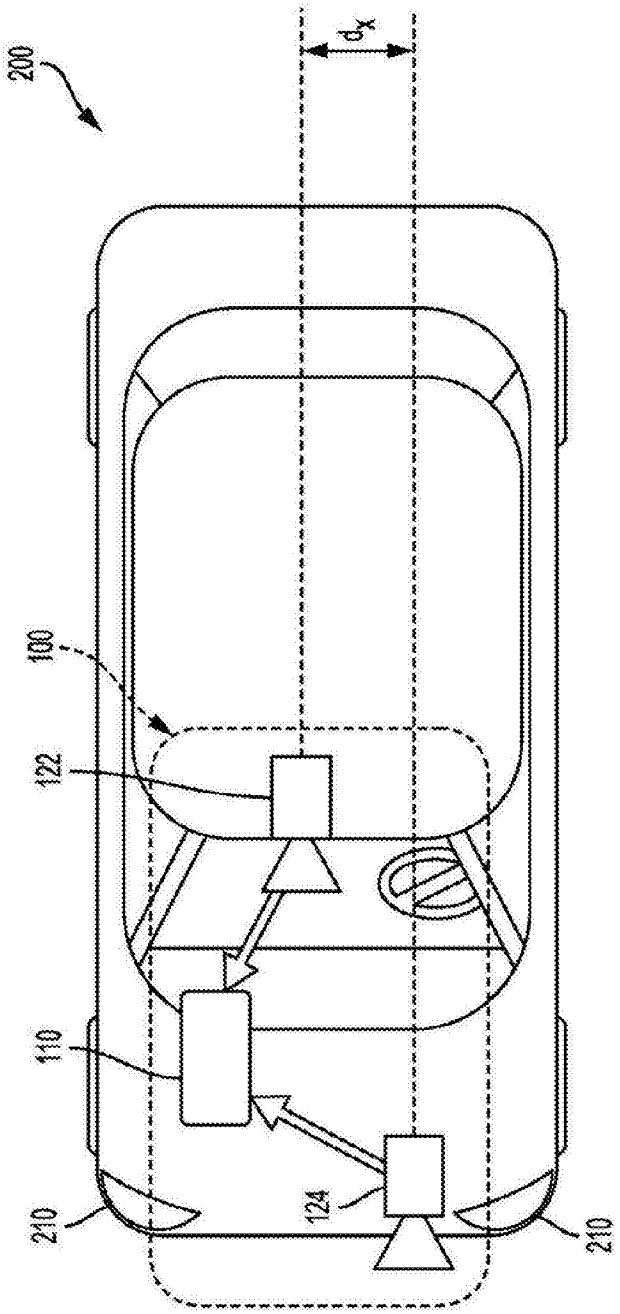


图2B

IPR2025-00943
Tesla EX1002 Page 1135

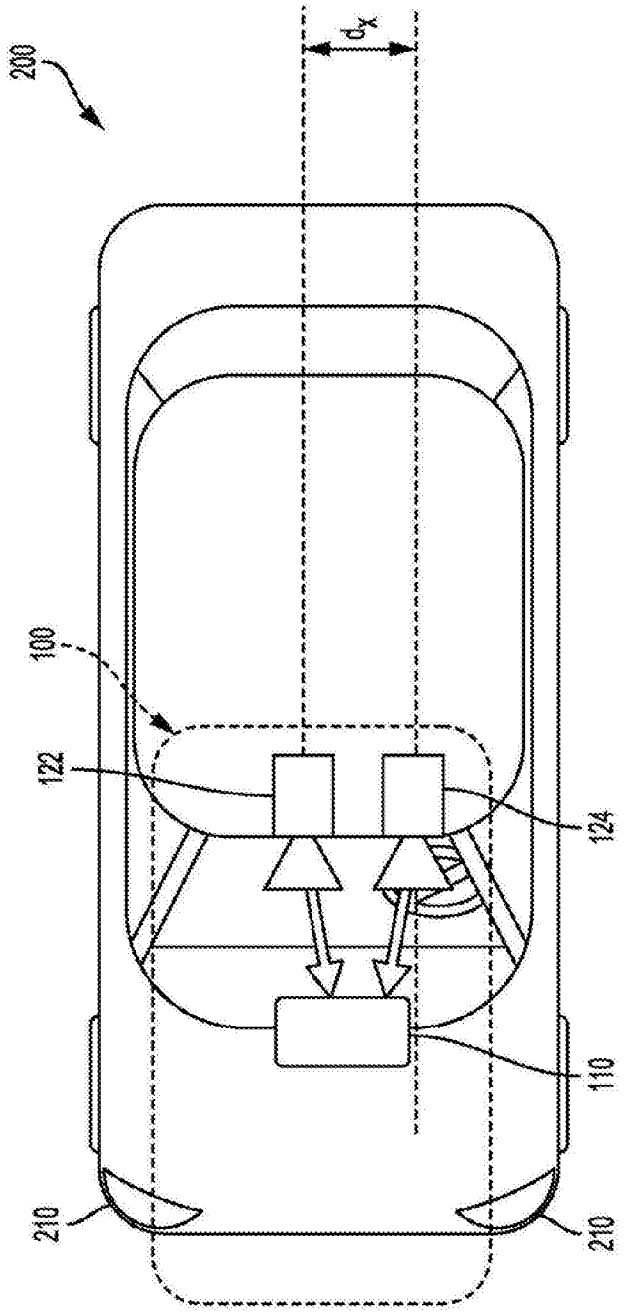


图2C

IPR2025-00943
Tesla EX1002 Page 1136

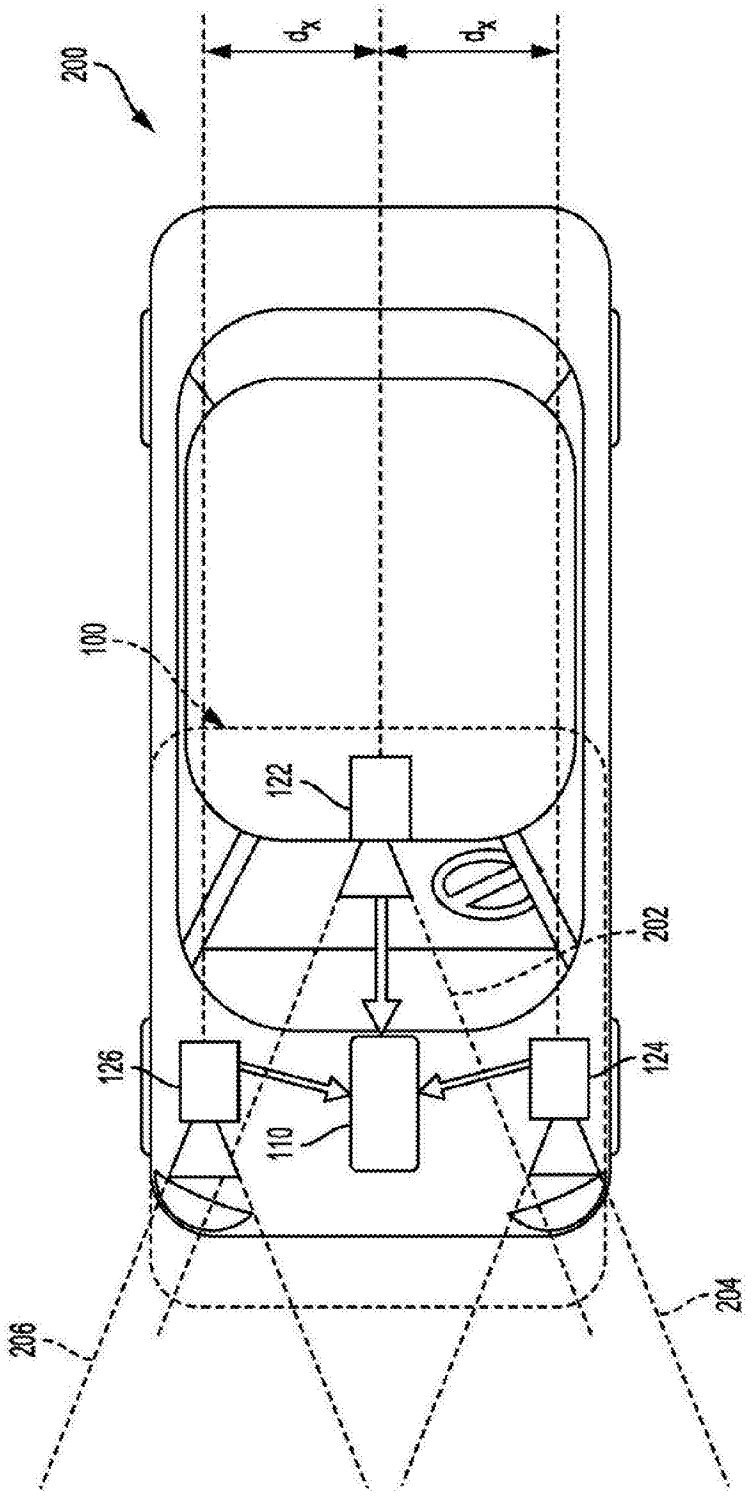


图2D

IPR2025-00943
Tesla EX1002 Page 1137

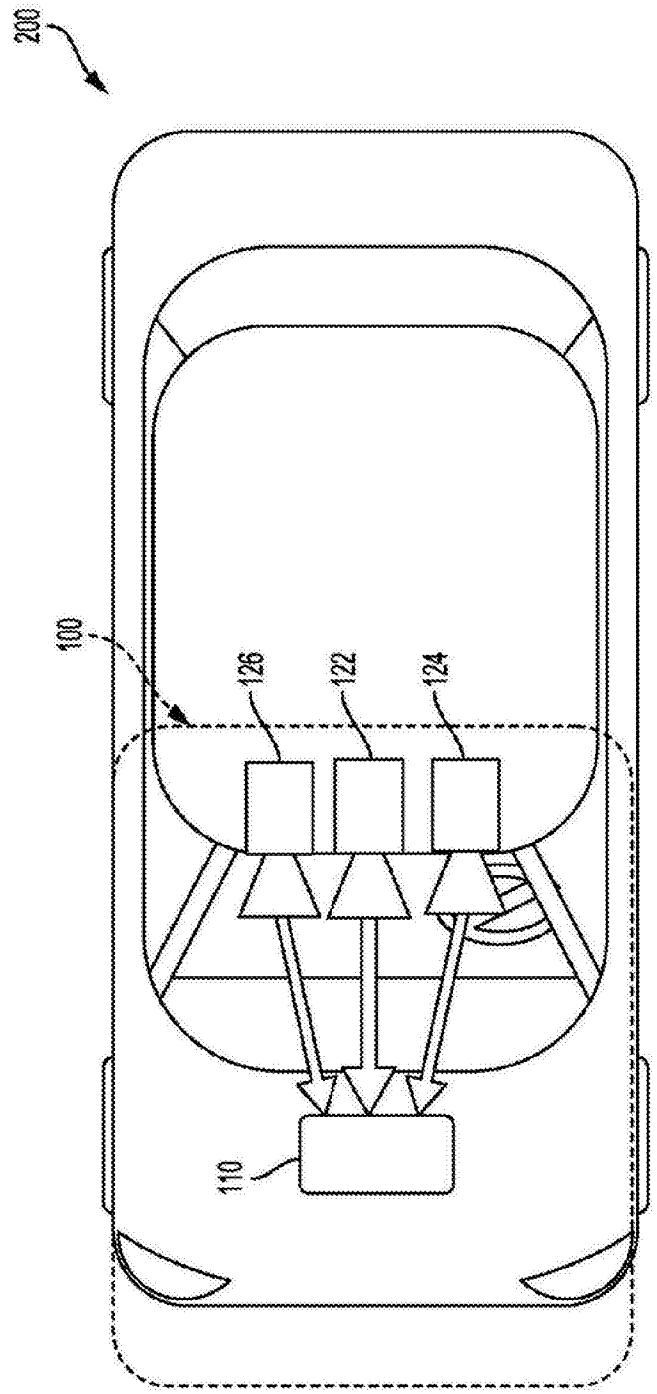


图2E

IPR2025-00943
Tesla EX1002 Page 1138

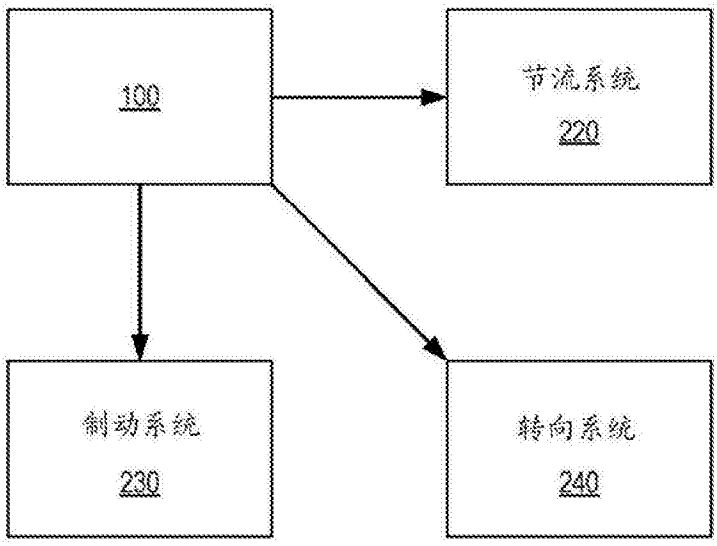
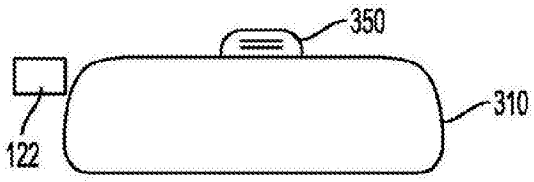


图2F

IPR2025-00943
Tesla EX1002 Page 1139



170

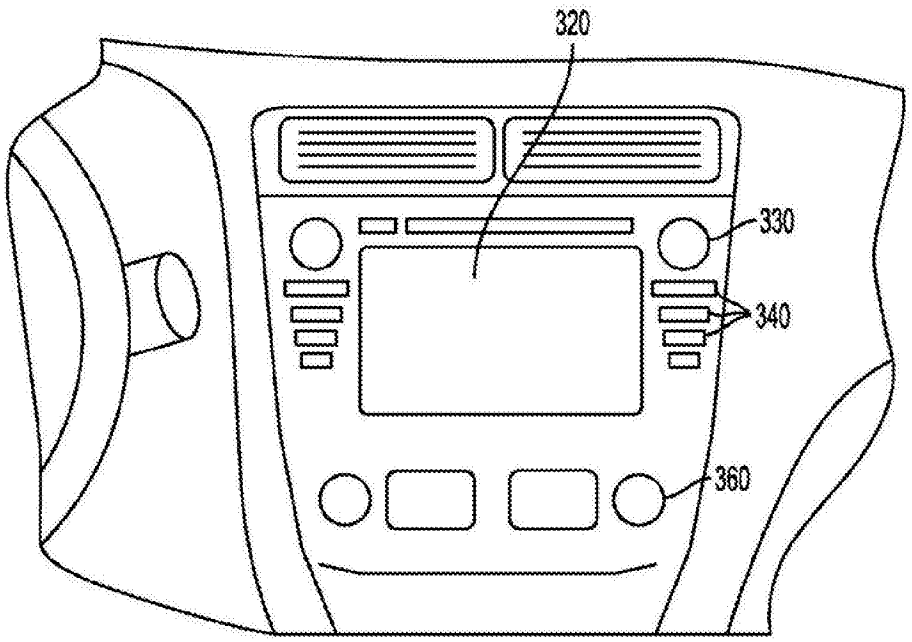


图3A

IPR2025-00943
Tesla EX1002 Page 1140

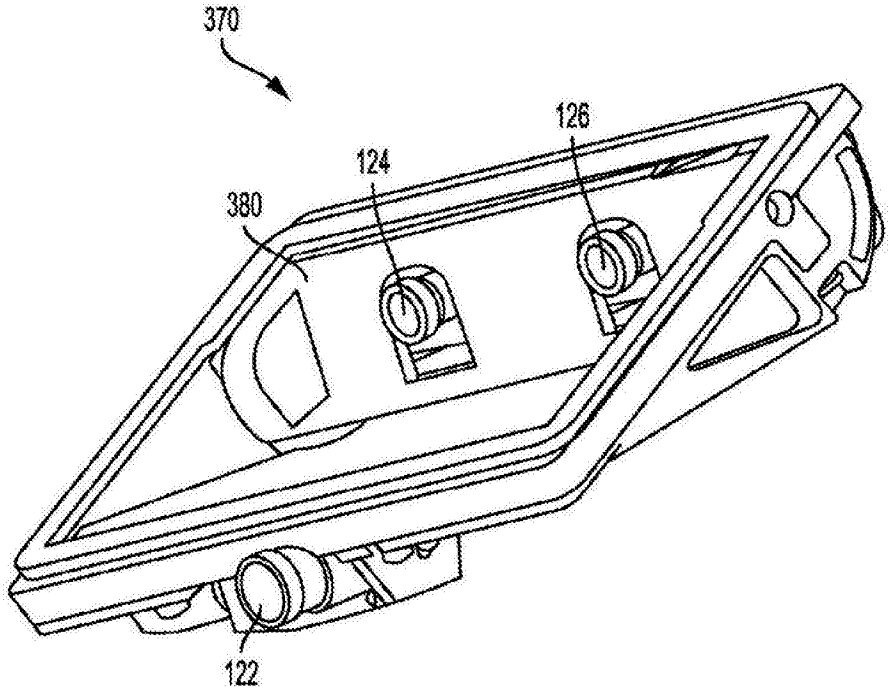


图3B

IPR2025-00943
Tesla EX1002 Page 1141

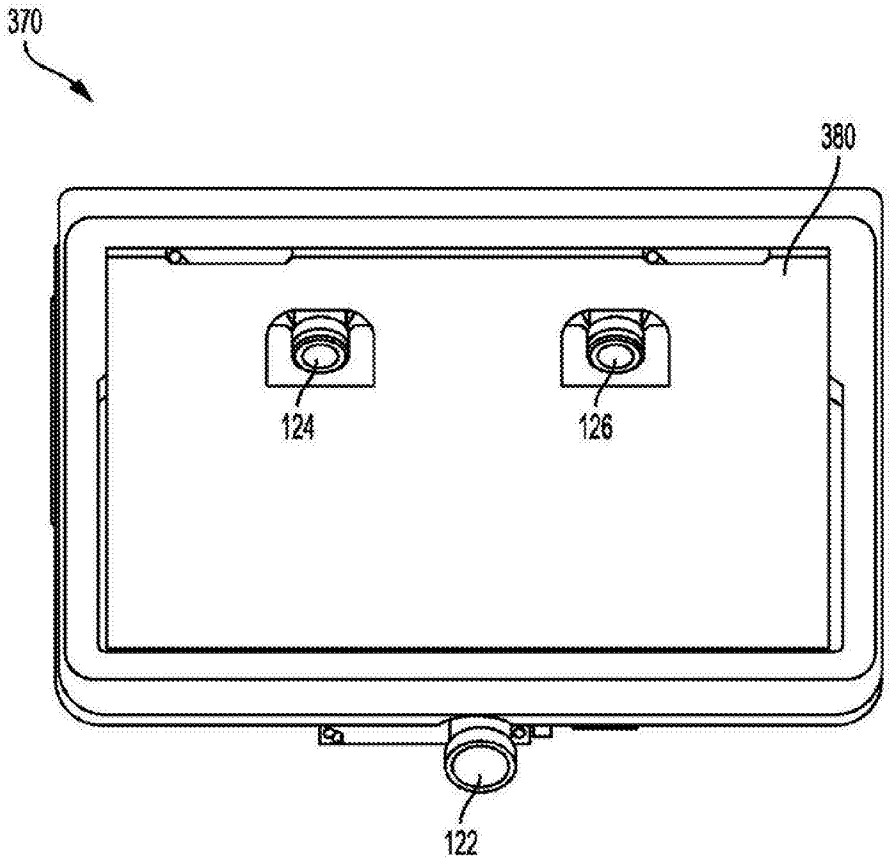


图3C

IPR2025-00943
Tesla EX1002 Page 1142

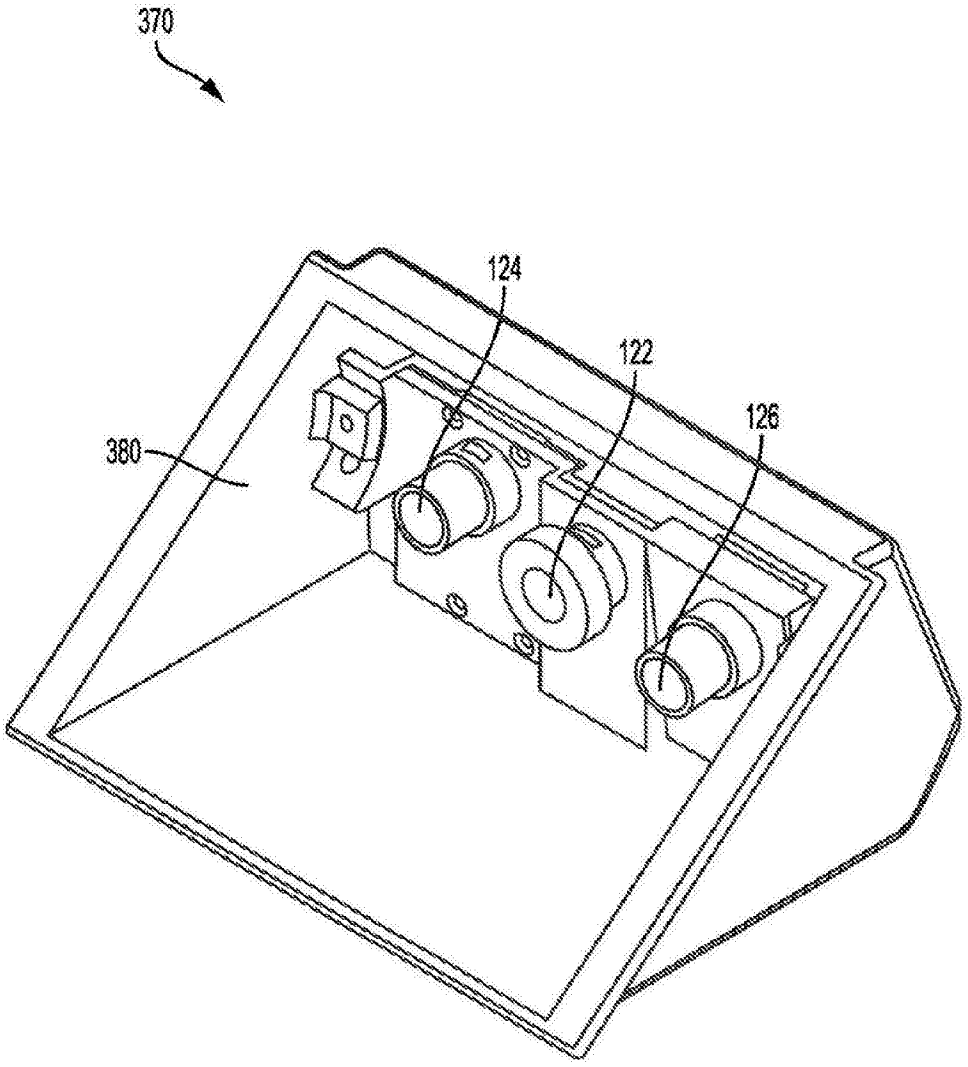


图3D

IPR2025-00943
Tesla EX1002 Page 1143

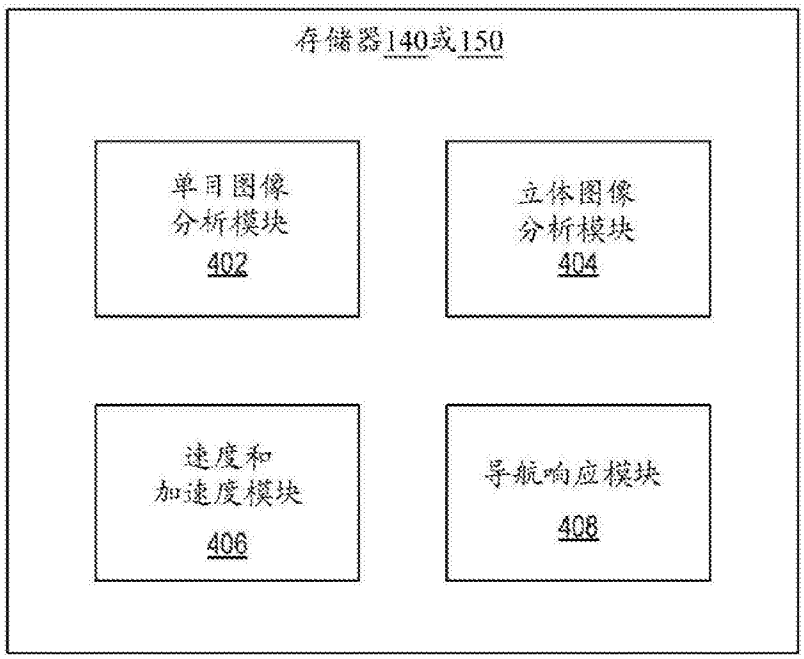


图4

IPR2025-00943
Tesla EX1002 Page 1144

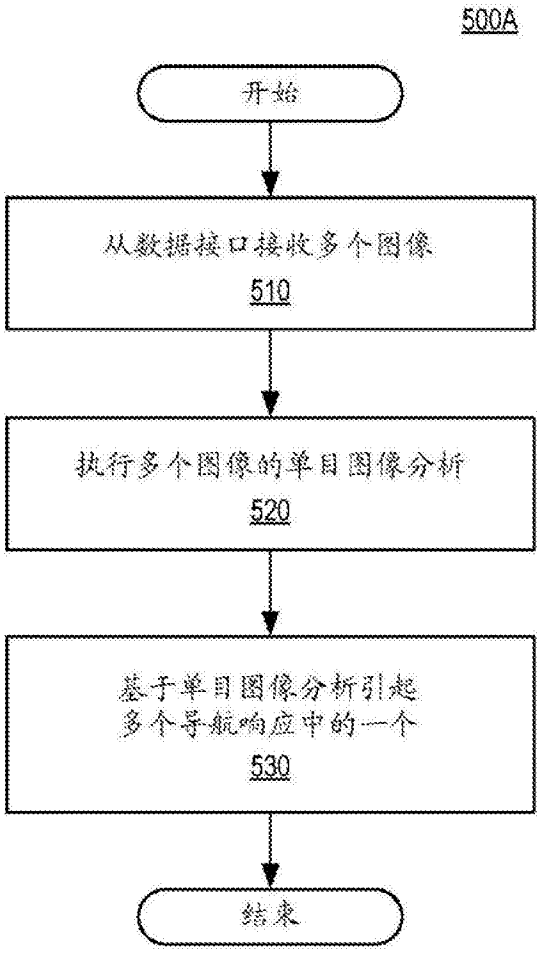


图5A

IPR2025-00943
Tesla EX1002 Page 1145

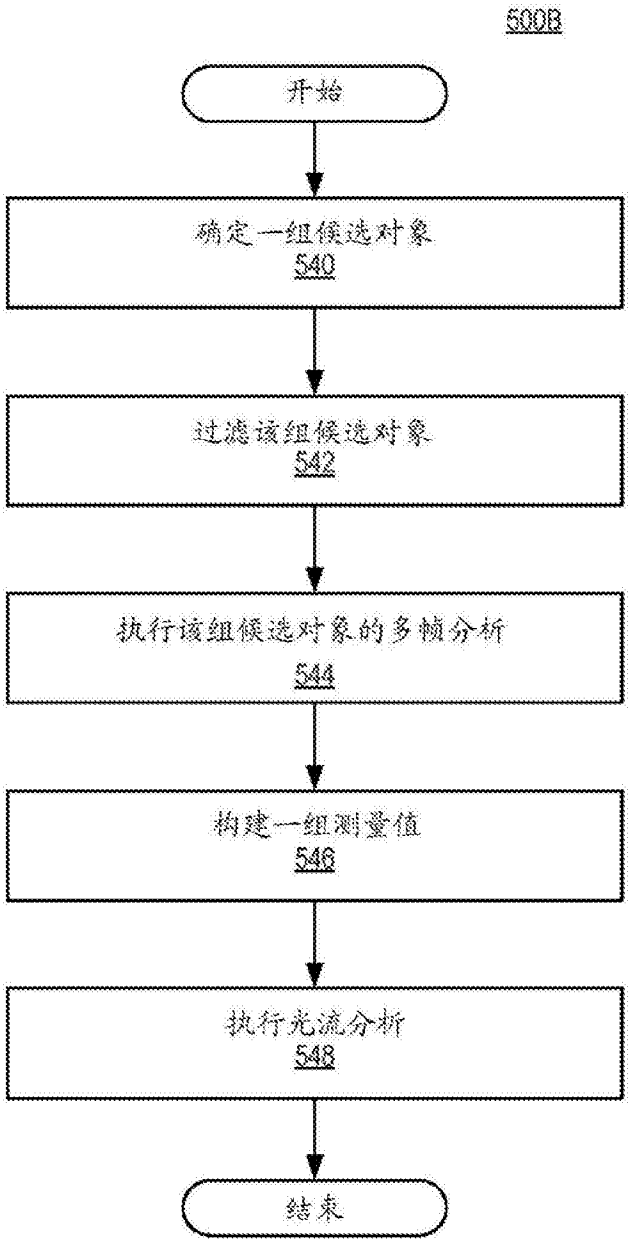


图5B

IPR2025-00943
Tesla EX1002 Page 1146

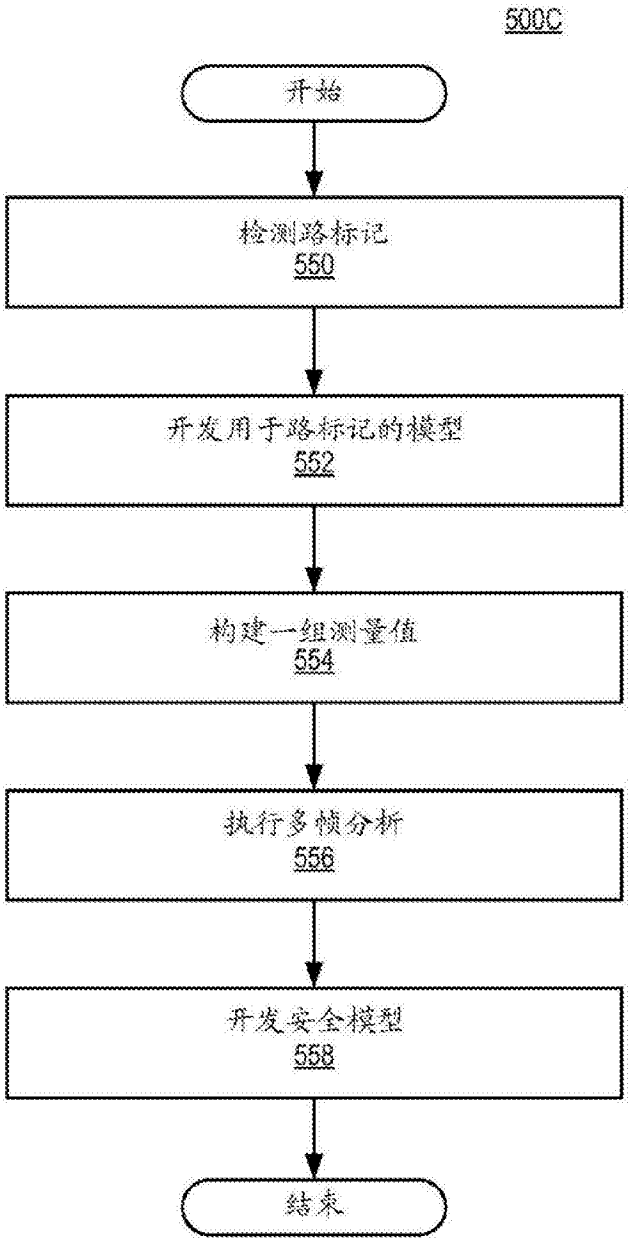


图5C

IPR2025-00943
Tesla EX1002 Page 1147

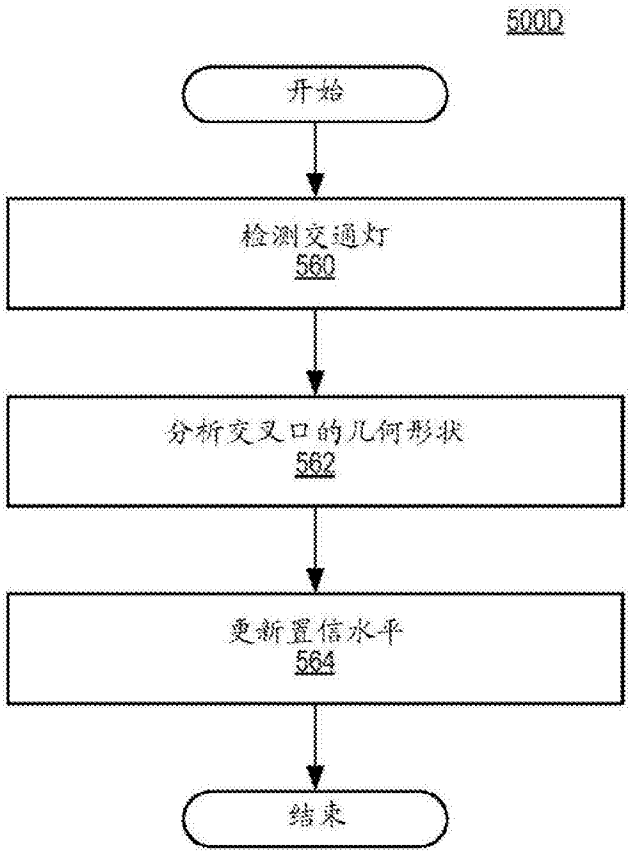


图5D

IPR2025-00943
Tesla EX1002 Page 1148

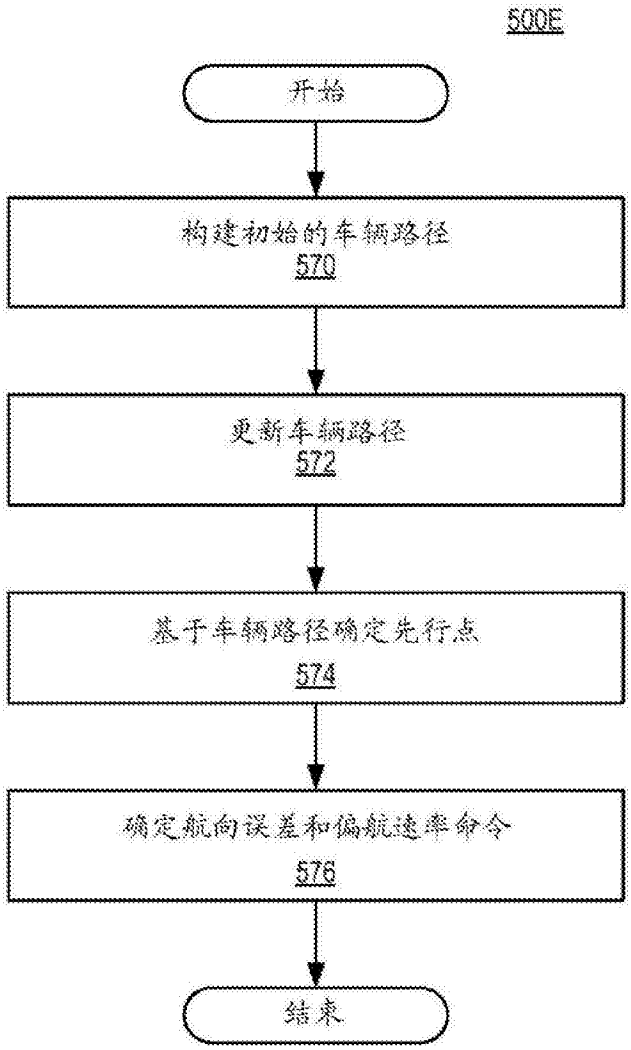


图5E

IPR2025-00943
Tesla EX1002 Page 1149

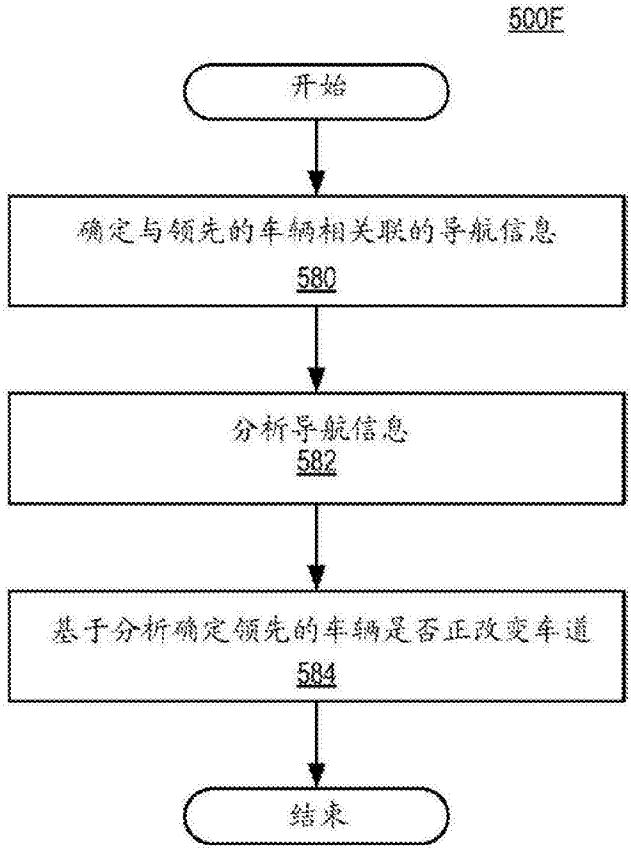


图5F

IPR2025-00943
Tesla EX1002 Page 1150

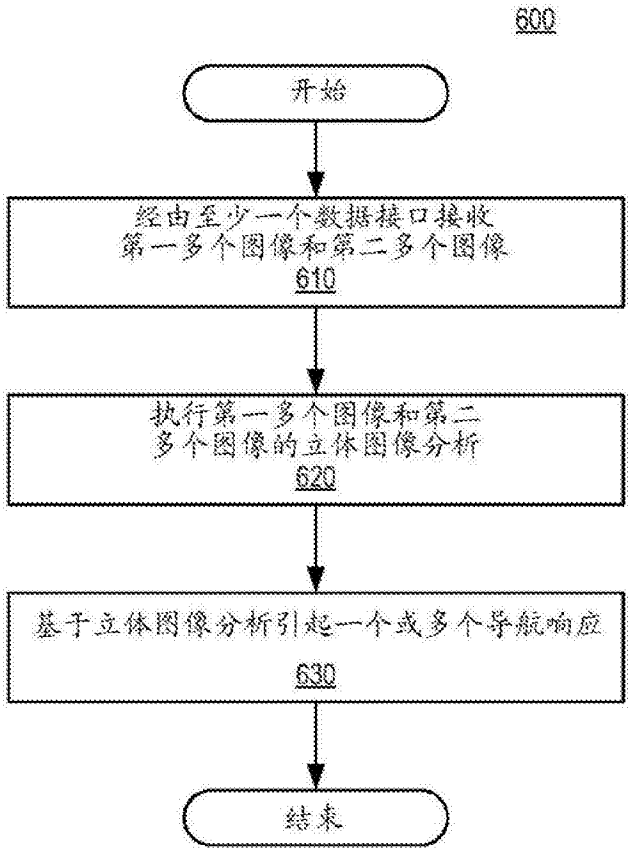


图6

IPR2025-00943
Tesla EX1002 Page 1151

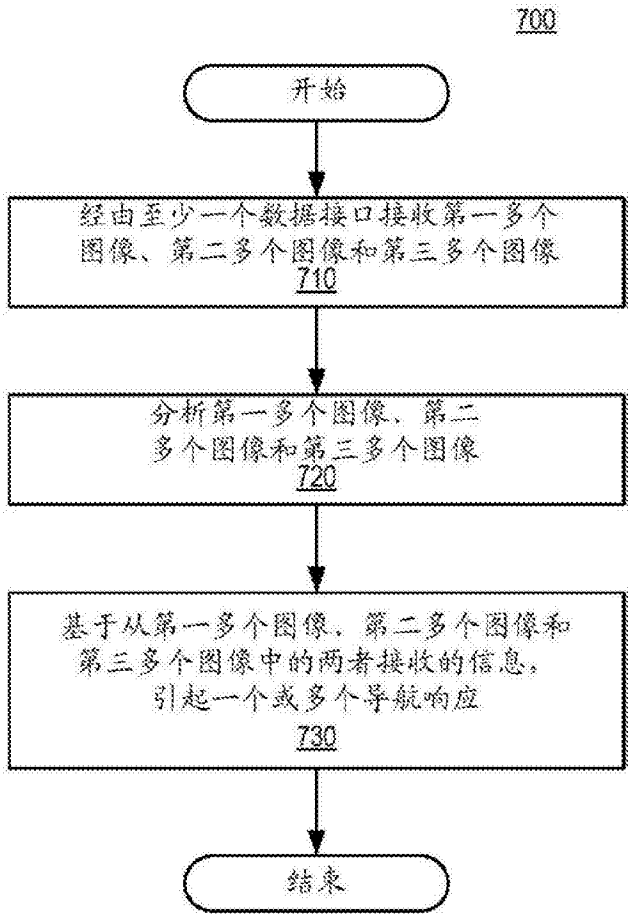


图7

IPR2025-00943
Tesla EX1002 Page 1152

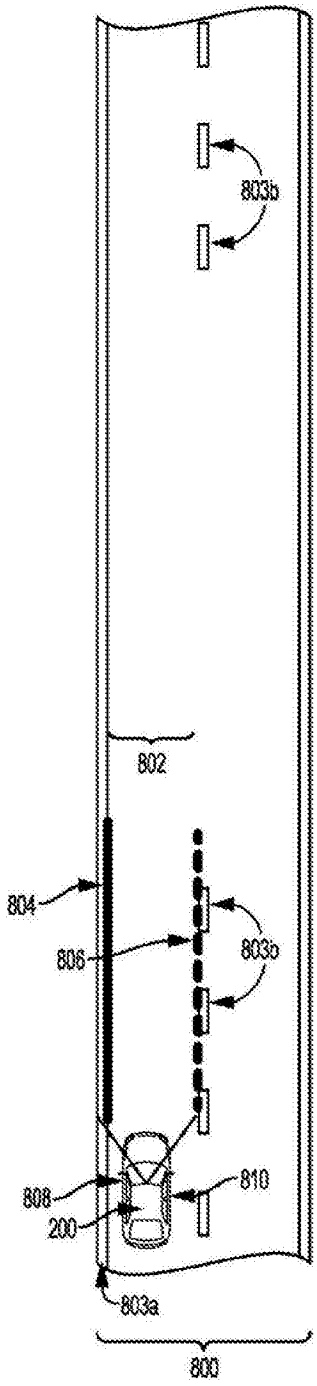


图8A

IPR2025-00943
Tesla EX1002 Page 1153

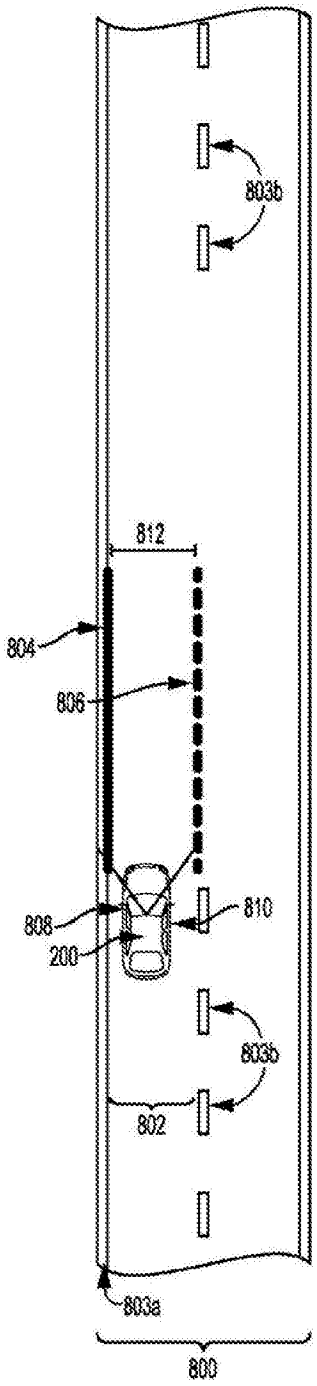


图8B

IPR2025-00943
Tesla EX1002 Page 1154

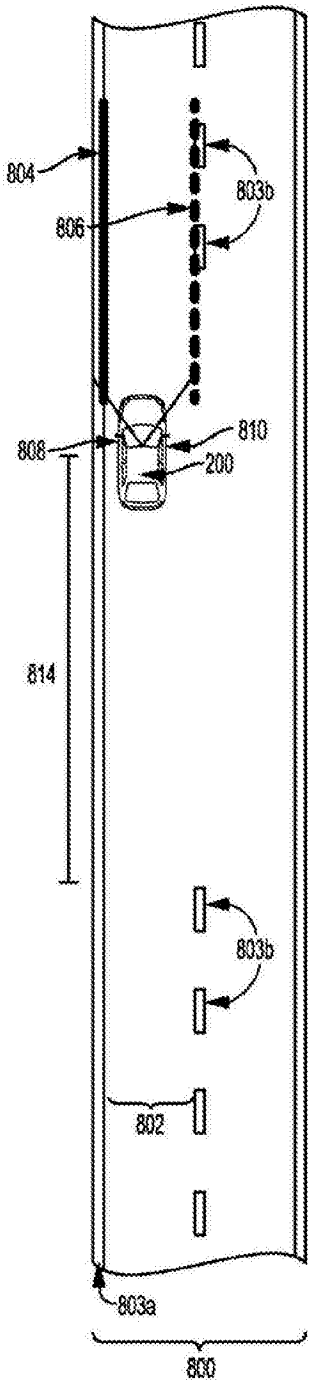


图8C

IPR2025-00943
Tesla EX1002 Page 1155

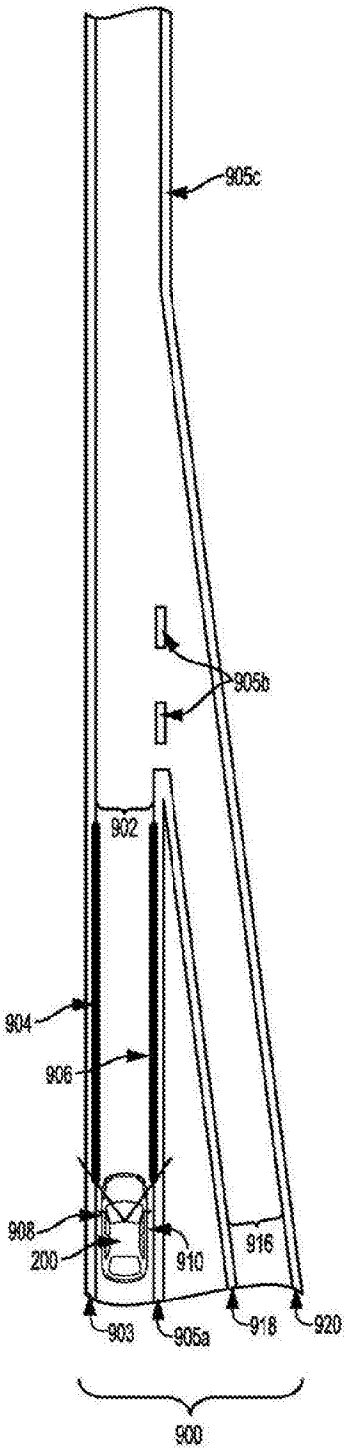


图9A

IPR2025-00943
Tesla EX1002 Page 1156

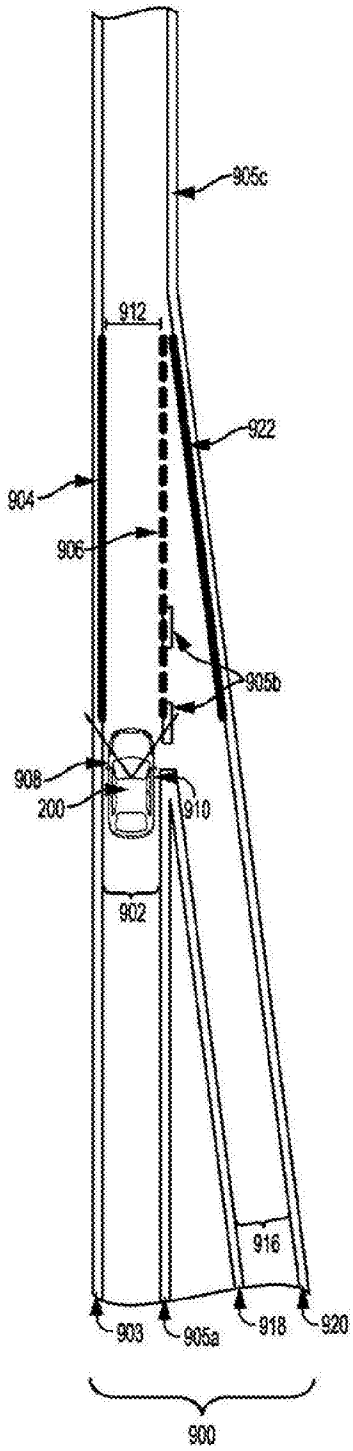


图9B

IPR2025-00943
Tesla EX1002 Page 1157

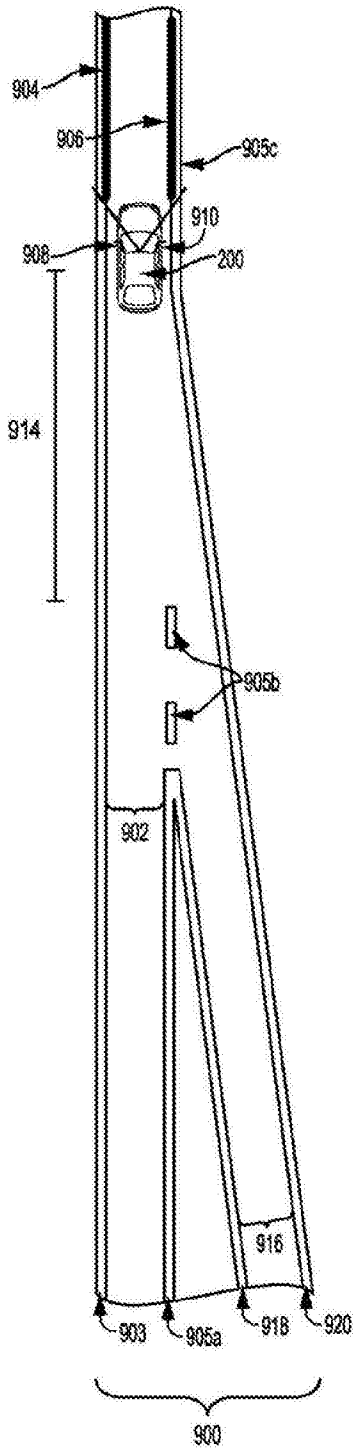


图9C

IPR2025-00943
Tesla EX1002 Page 1158

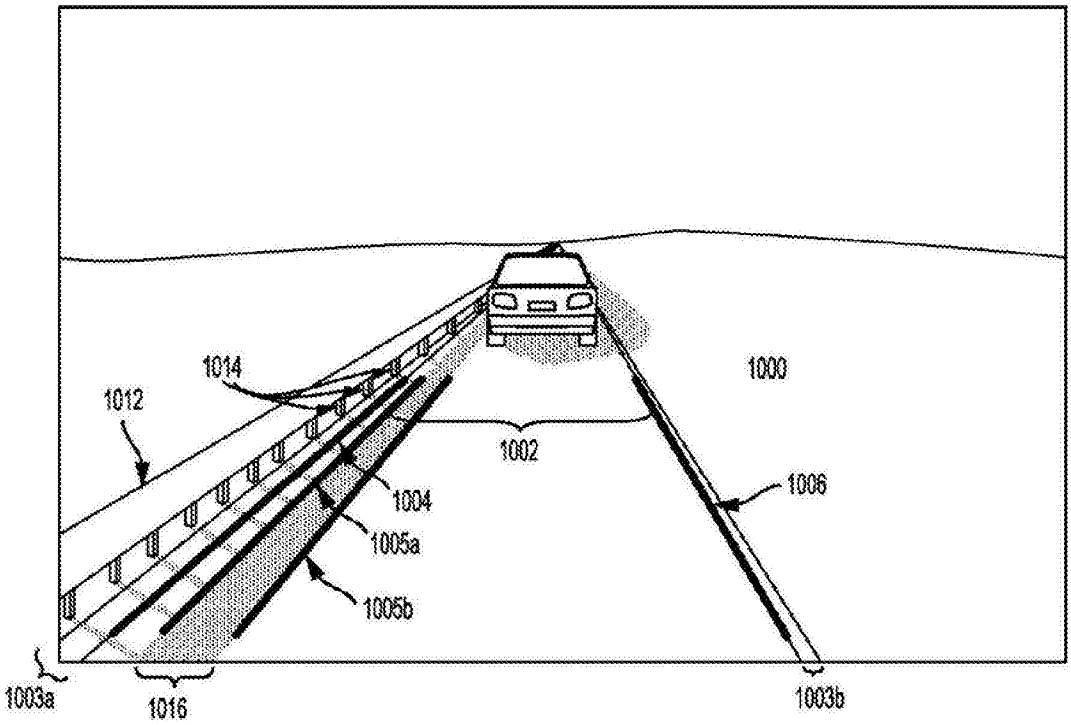


图10A

IPR2025-00943
Tesla EX1002 Page 1159

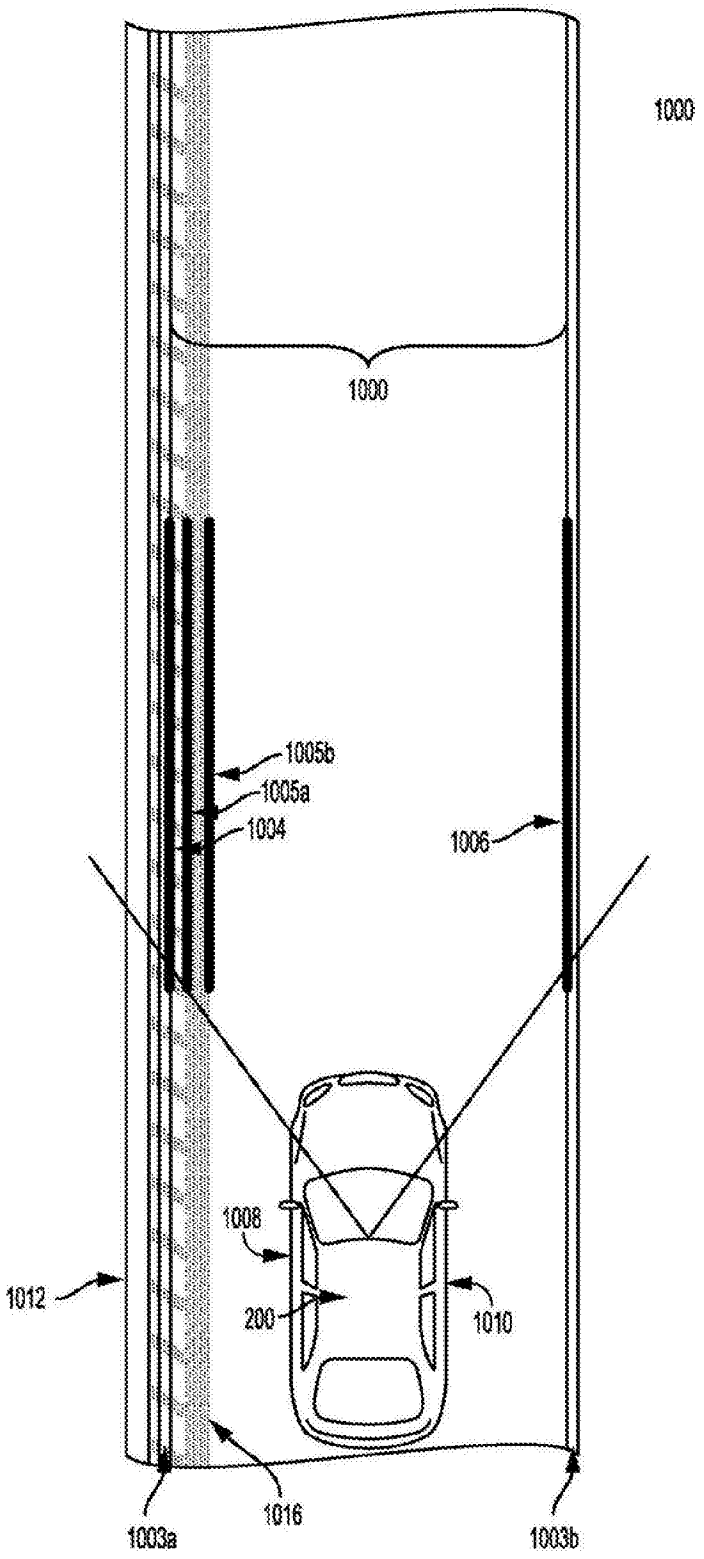


图10B

IPR2025-00943
Tesla EX1002 Page 1160

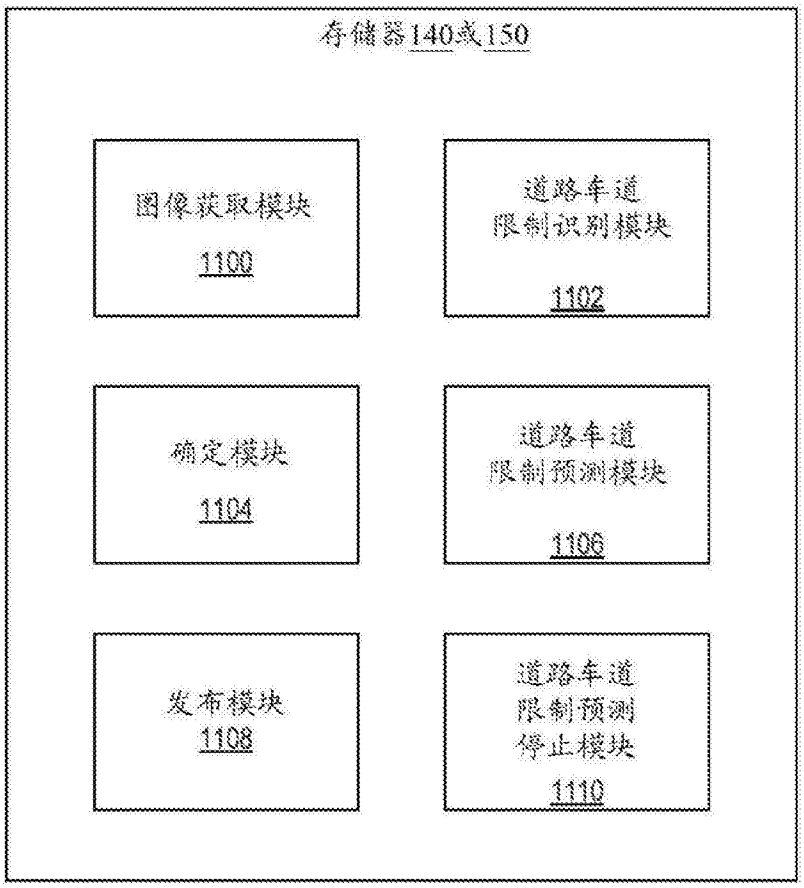


图11

IPR2025-00943
Tesla EX1002 Page 1161

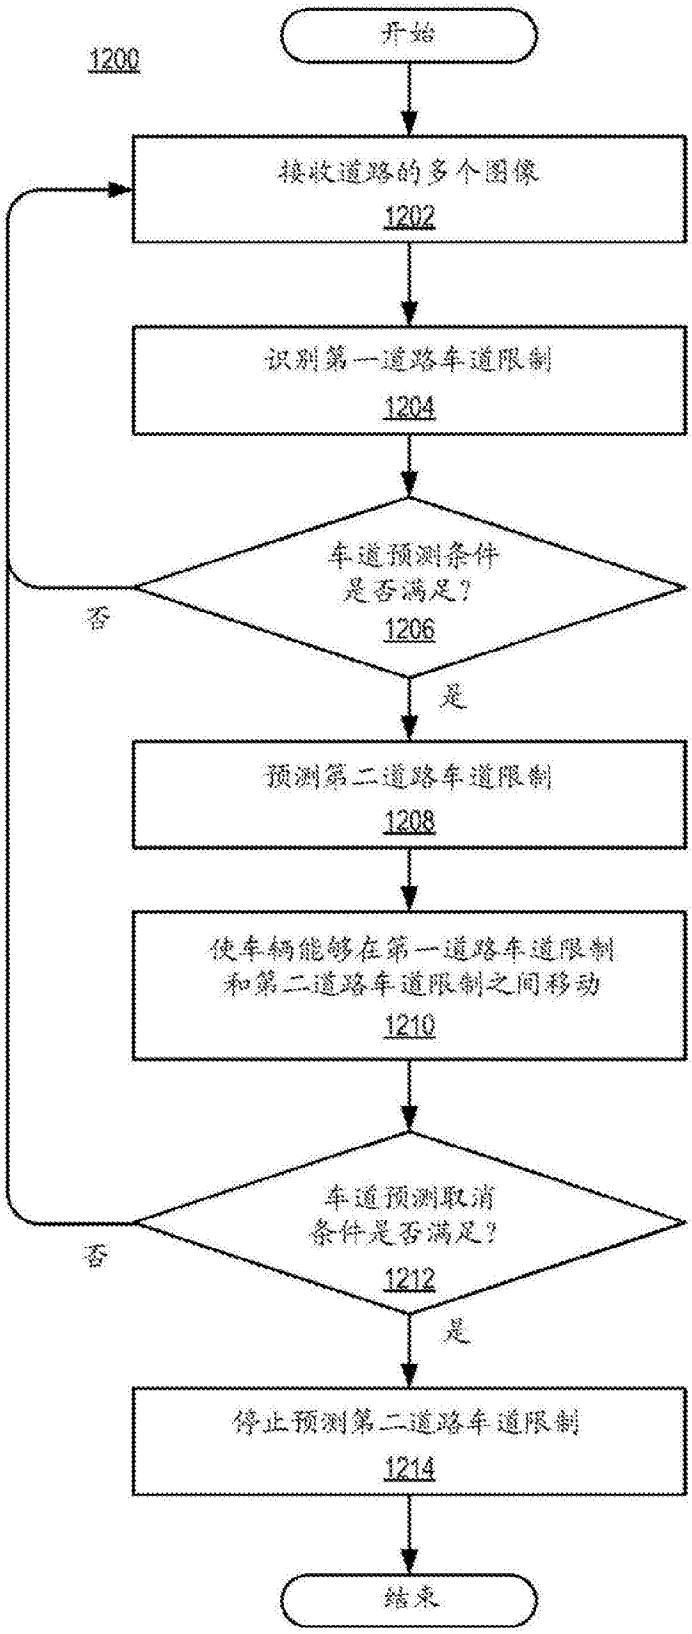


图12

IPR2025-00943
Tesla EX1002 Page 1162

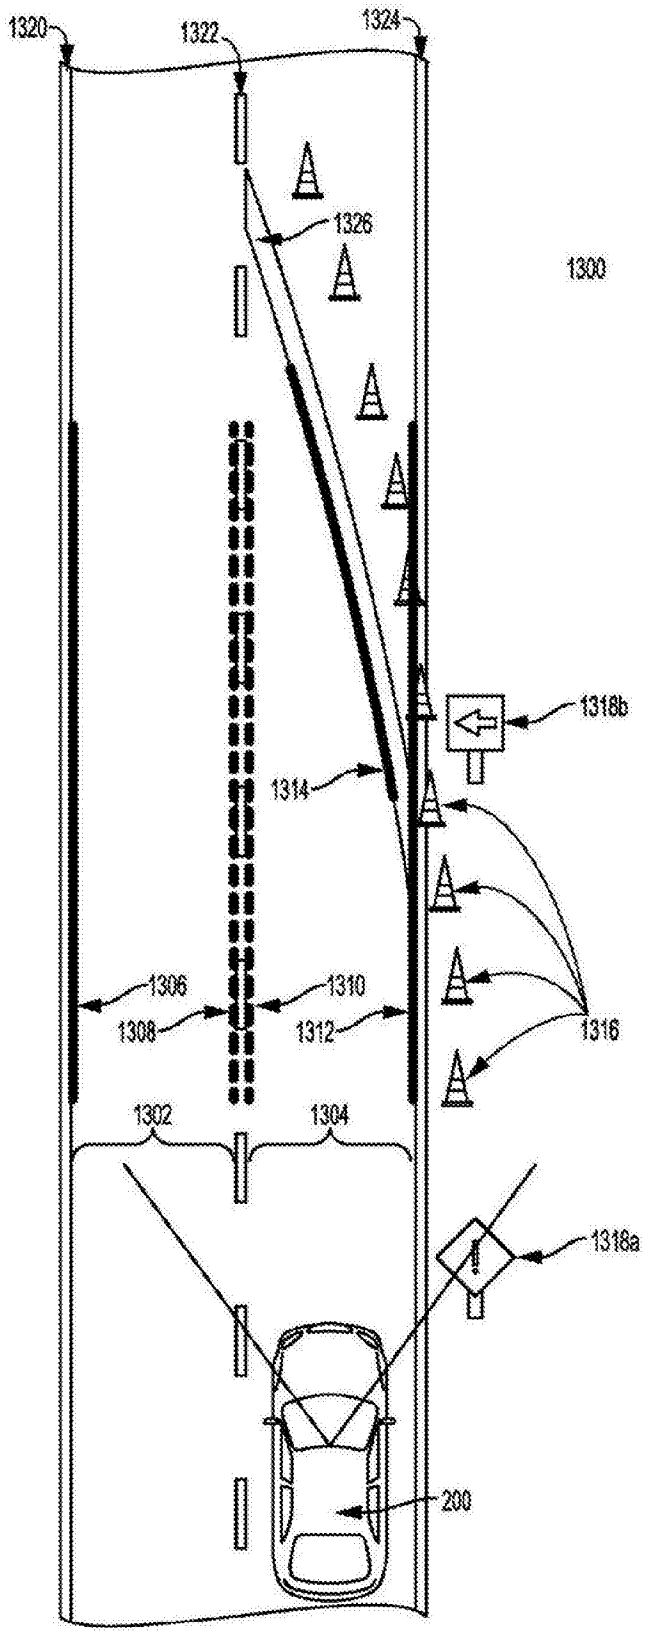


图13

IPR2025-00943
Tesla EX1002 Page 1163

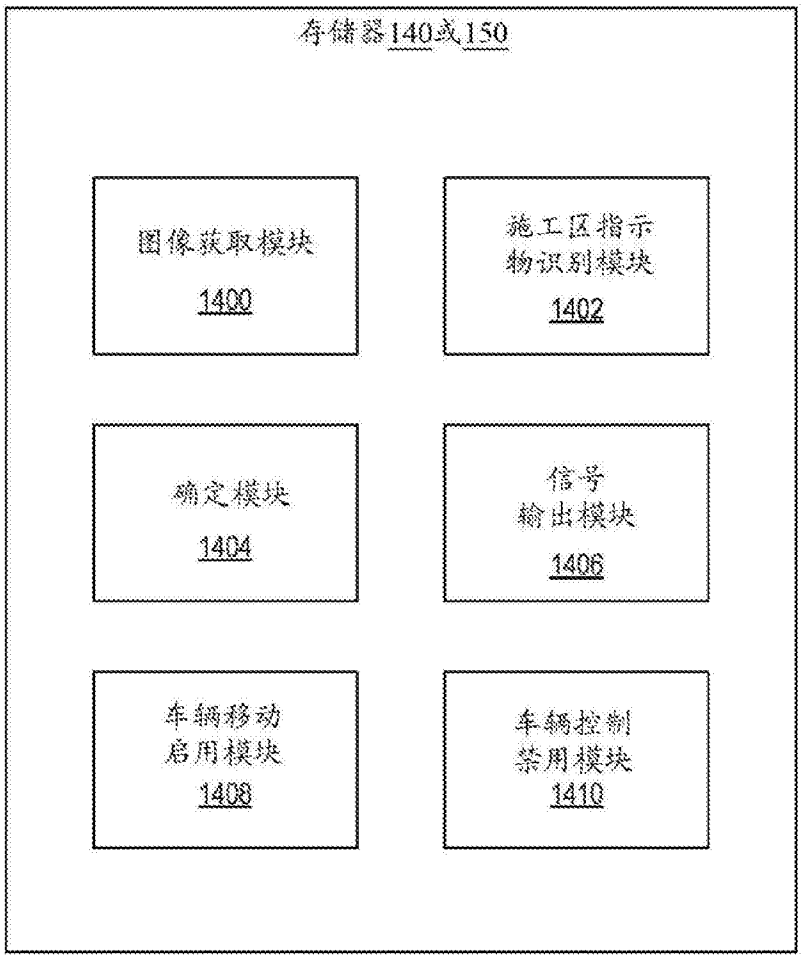


图14

IPR2025-00943
Tesla EX1002 Page 1164

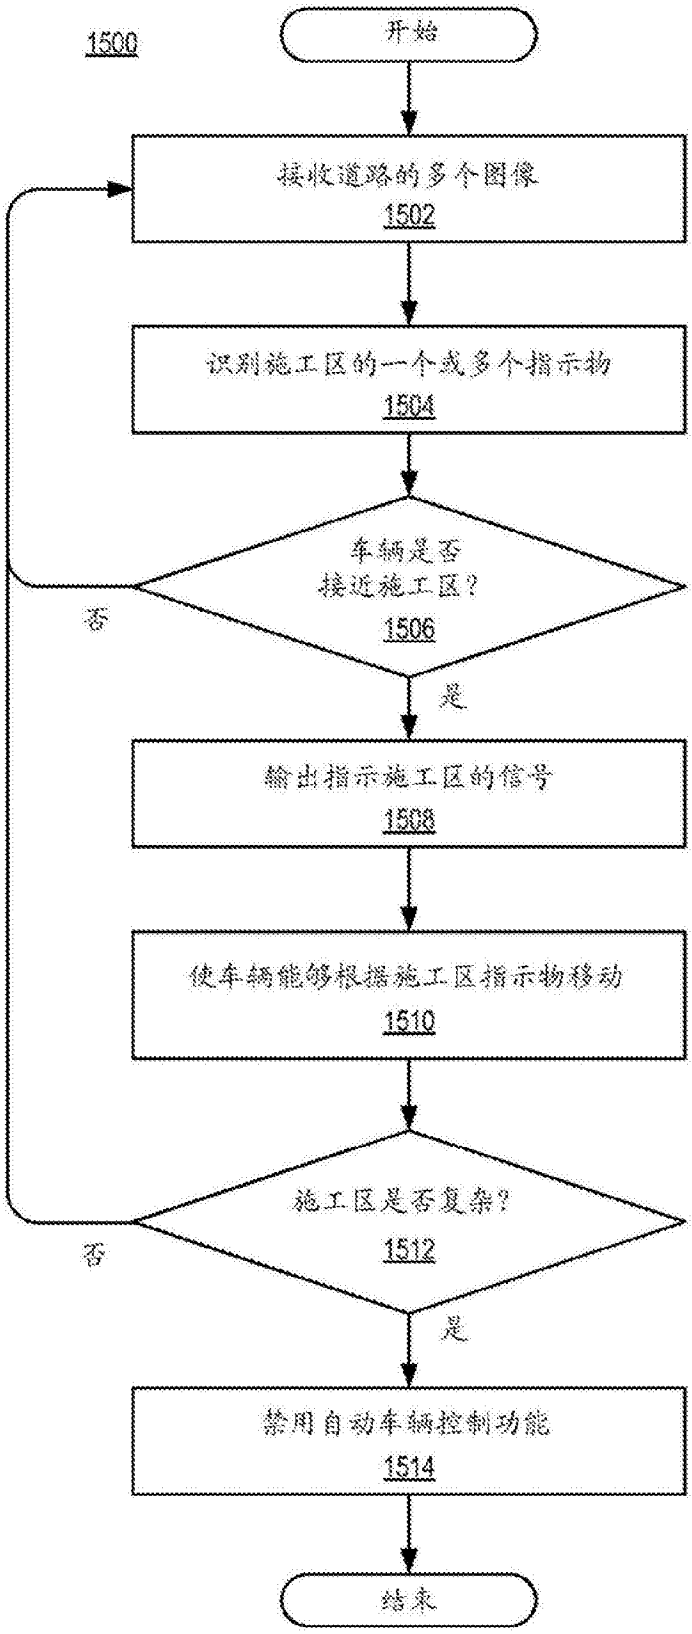


图15

IPR2025-00943
Tesla EX1002 Page 1165

Espacenet

Bibliographic data: CN110920609 (A) — 2020-03-27

Systems and methods for mimicking leading vehicle

Inventor(s): CLARKE ANNA; BAGON EYA; NIR GALI ± (A·克拉克, ; E·宝葛, ; G·尼尔)

Applicant(s): MOBILEYE VISION TECHNOLOGIES LTD ± (移动眼视力科技有限公司)

Classification:

- **international:** ***B60W30/12; B60W30/18; G01C21/34; G01C21/36; G06K9/00***
- **cooperative:** **B60K31/0066 (US)**; **B60R1/00 (US)**; **B60T13/662 (EP)**; **B60T7/12 (US)**; **B60T7/22 (EP, US)**; **B60T8/32 (US)**; **B60T8/321 (EP)**; **B60W10/18 (US)**; **B60W10/20 (US)**; **B60W30/00 (US)**; **B60W30/08 (US)**; **B60W30/09 (US)**; **B60W30/0953 (US)**; **B60W30/12 (US)**; **B60W30/14 (US)**; **B60W30/143 (US)**; **B60W30/165 (EP, US)**; **B60W30/18154 (EP, US)**; **B60W30/18163 (EP, US)**; **B60W40/06 (US)**; **B60W40/072 (US)**; **B60W40/076 (US)**; **B60W50/14 (US)**; **B62D15/025 (EP)**; **B62D15/0255 (EP)**; **B62D15/0265 (EP, US)**; **G01C11/04 (US)**; **G01C21/10 (US)**; **G01C21/26 (US)**; **G01C21/30 (EP, US)**; **G05D1/0088 (US)**; **G05D1/0246 (US)**; **G06K9/00798 (EP, US)**; **G06K9/00805 (EP)**; **G06K9/00825 (EP)**; **G06T7/285 (US)**; **G06T7/292 (US)**; **G06T7/70 (US)**; **G06T7/97 (US)**; **G08G1/09623 (EP, US)**; **G08G1/166 (EP, US)**; **G08G1/167 (EP, US)**; **H04N5/247 (US)**; B60K2031/0025 (US); B60T2201/022 (EP, US); B60T2201/03 (EP); B60T2201/08 (EP, US); B60T2210/34 (EP, US); B60W2050/0075 (EP, US); B60W2420/403 (US); B60W2420/42 (EP, US); B60W2510/0604 (US); B60W2510/18 (US); B60W2510/20 (US); B60W2520/14 (US); B60W2552/30 (EP); B60W2554/4041 (EP, US); B60W2554/80 (US); B60W2556/10 (EP, US); B60W2556/50 (EP, US); B60W2710/18 (US); B60W2710/20 (US); B62D15/0255 (US); B62D6/00 (US); G01S5/0027 (US); G05D1/0214 (US); G05D1/0253 (US); G05D1/0278 (US); G06T2207/10028 (US); G06T2207/30256 (US); G06T2207/30261 (US); G08G1/09 (US); G08G1/22 (EP, US); H04N7/18 (US)

IPR2025-00943
Tesla EX1002 Page 1166

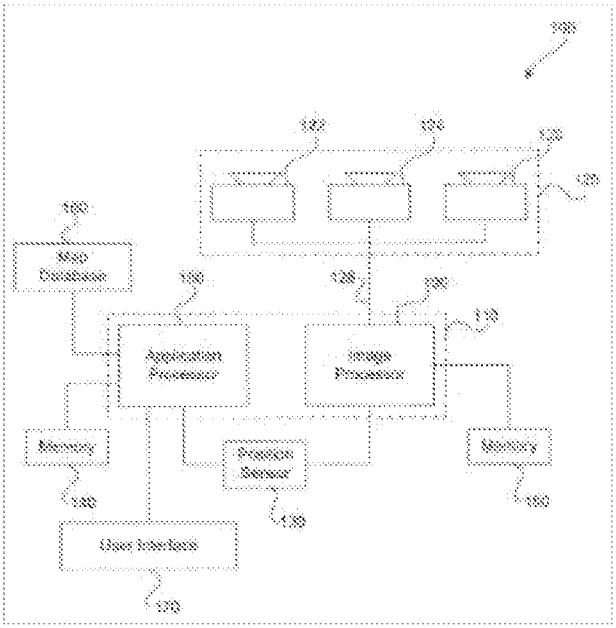
Application number: CN201911164109 20141204 Global Dossier

Priority number(s): US201361911490P 20131204 ; US201461993084P 20140514 ; US201461993111P 20140514 ; US201462015524P 20140623 ; US201462022221P 20140709 ; US201462040224P 20140821 ; US201462040269P 20140821 ; CN20148074153 20141204

Also published as: CN106030609 (A) CN106030609 (B) CN111027420 (A) EP3092599 (A1) EP3092599 (B1) more

Abstract of CN110920609 (A)

The invention relates to systems and methods for mimicking a leading vehicle. Systems and methods use cameras to provide autonomous navigation features. In one implementation, a driver assist navigation system is provided for a primary vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the primary vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: locate in the plurality of images a leading vehicle; determine, based on the plurality of images, at least one action taken by the leading vehicle; and cause the primary vehicle to mimic the at least one action of the leading vehicle.



Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

DESCRIPTION CN110920609A

System and method for imitating the preceding vehicle

[0001]

This application is a divisional application for a patent application whose filing date is December 4, 2014, the application number is 201480074153.9 (PCT/IB2014/003151), and the invention title is "System and method for imitating the preceding vehicle".

[0002]

Technical field

[0003]

The present disclosure generally relates to autonomous vehicle navigation, and more specifically, to systems and methods that use cameras to provide autonomous vehicle navigation features.

[0004]

Background technique

[0005]

With the continuous advancement of technology, the goal of a fully autonomous vehicle capable of navigating on the road is about to be reached.

First, the autonomous vehicle can recognize its environment and navigate without operator input.

Autonomous vehicles can also consider various factors and make appropriate decisions based on these factors to safely and accurately reach the intended destination.

For example, when a vehicle habitually travels on a road, it encounters various objects such as other vehicles and pedestrians.

The autonomous driving system can recognize these objects in the vehicle environment and take appropriate and timely actions to avoid collisions. In addition, the autonomous drive system can recognize other indicators, such as traffic signals, traffic signs, and lane markings, which regulate vehicle movement (for example, when the vehicle must stop and can drive, the vehicle must not exceed the speed, the vehicle must be positioned On the road, etc.). Autonomous drive systems may need to determine when the vehicle should change lanes, turn at intersections, change roads, and so on. It is obvious from these examples that many factors may need to be addressed in order to provide an autonomous vehicle that can navigate safely and accurately.

[0006]

Summary of the invention

[0007]

Embodiments consistent with the present disclosure provide systems and methods for autonomous vehicle navigation.

The disclosed embodiments may use cameras to provide autonomous vehicle navigation features.

For example, consistent with the disclosed embodiments, the disclosed system may include one, two or more cameras for monitoring the vehicle environment and generating based on the analysis of the images captured by one or more of the cameras. Navigation response.

[0008]

Consistent with the disclosed embodiments, a driver assistance navigation system for a vehicle is provided.

The system may include: at least one image capturing device configured to acquire a plurality of images of an area near the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine the first lane constraint on the first side of the vehicle based on the plurality of images; Determine a second lane constraint on a second side of the vehicle opposite to the first side of the vehicle, wherein the first lane constraint and the second lane constraint define the lane in which the vehicle travels, and wherein The first distance corresponds to the distance between the first side of the vehicle and the first lane constraint, and the second distance corresponds to the distance between the second side of the vehicle and the second lane constraint的 distance; based on the multiple images to determine whether the lane offset condition exists on the first side of the vehicle; if the lane offset condition exists on the first side of the vehicle, cause all The vehicle travels within the first lane constraint and the second lane constraint, so that the first distance is greater than the second distance; based on the multiple images, it is determined whether a lane offset condition exists in all of the vehicle. And if a lane offset condition exists on the second side of the vehicle, then cause the vehicle to travel within the first and second lane constraints, so that the first The distance is smaller than the second distance.

[0009]

Consistent with the disclosed embodiment, the at least one processing device may be further configured to, if it is determined based on the plurality of images that a lane shift condition exists on both the first side and the second side of the vehicle, then As a result, the first distance and the second distance are substantially the same.

The at least one processing device may be configured to determine, based on the plurality of images, whether a lane deviation condition is substantially non-existent, and if the lane deviation condition is substantially non-existent, cause the vehicle to be in the first The lane constraint and the second lane constraint travel within such that the first distance is substantially equal to the second distance. If the driver assistance navigation system determines that a lane deviation condition exists, the at least one processing device may be configured to cause the first distance to be at least 1.25 times greater than the second distance. Alternatively, the at least one processing device may be further configured to cause the first distance to be at least 2 times greater than the second distance. If the driver assistance navigation system determines that the lane deviation condition includes a curved road, the at least one processing device may be configured to cause the vehicle to travel within the first lane constraint and the second lane constraint, so that the The first distance is toward the outside of the curve relative to the vehicle, and the second distance is toward the inside of the curve relative to the vehicle. If the driver assistance navigation system determines that the lane offset condition includes the presence of one or more objects on one side of the lane, the at least one processing device may be configured to cause the vehicle to be in the first The lane constraint and the second lane constraint travel within such that the first distance greater than the second distance is provided toward the one or more objects. The one or more objects may include parked cars, walls, and/or pedestrians, and may have a height greater than 10 cm from the road surface. At least one of the first lane constraint and the second lane constraint may include a line marked on the road surface, a curb of the road surface, or a determined midpoint of a road surface width. Determining whether a lane deviation condition exists on the first side of the vehicle may be performed based on the position of the target object on the first side of the vehicle. Causing the vehicle to travel within the first lane constraint and the second lane constraint such that the first distance is greater than the second distance may be revoked based on the input of the driver of the vehicle.

[0010]

Tesla EX1002 Page 1169 IPR2025-00943

IPR2025-00943 Tesla EX1002 Page 1170

Consistent with another disclosed embodiment, a vehicle may include: a first vehicle side; a second vehicle side opposite to the first vehicle side; at least one image capturing device, the at least one image capturing device being configured To obtain multiple images of the area near the vehicle; a data interface; and at least one processing device.

The at least one processing device may be configured to: receive the plurality of images via the data interface; determine a first lane constraint on the first vehicle side based on the plurality of images; determine based on the plurality of images A second lane constraint on the second vehicle side, wherein the first lane constraint and the second lane constraint define the lane in which the vehicle travels, and wherein the first distance corresponds to the first vehicle side and the The distance between the first lane constraint, and the second distance corresponds to the distance between the second vehicle side and the second lane constraint; based on the plurality of images, it is determined whether a lane offset condition exists in the On the first vehicle side; if the lane offset condition exists on the first vehicle side, the vehicle is caused to travel within the first lane constraint and the second lane constraint, so that the first distance is greater than the Second distance; based on the multiple images to determine whether the lane offset condition exists on the second vehicle side; and if the lane offset condition exists on the second vehicle side, cause the vehicle to be at all The first lane constraint and the second lane constraint travel within such that the first distance is smaller than the second distance.

[0011]

Consistent with the disclosed embodiment, the at least one processing device may be further configured to cause all lane offset conditions to exist on both the first vehicle side and the second vehicle side based on the plurality of images. The first distance is substantially the same as the second distance.

The at least one processing device may be configured to determine, based on the plurality of images, whether a lane deviation condition is substantially non-existent, and if the lane deviation condition is substantially non-existent, cause the vehicle to be in the first The lane constraint and the second lane constraint travel within such that the first distance is substantially equal to the second distance. The lane deviation condition may include a curved road, and the at least one processing device may be configured to cause the vehicle to travel within the first lane constraint and the second lane constraint such that the first distance faces the opposite of the curve. On the outside of the vehicle, and the second distance is toward the inside of the curve relative to the vehicle. The lane deviation condition may include the presence of one or more objects on one side of the lane, and the at least one processing device may be configured to cause the vehicle to be within the first lane constraint and the second lane constraint Travel such that the first distance greater than the second distance is provided toward the one or more objects. At least one of the first lane constraint and the second lane constraint may include a curb of the road surface and/or a determined midpoint of the road surface width. Determining whether a lane deviation condition exists on the first side of the vehicle may be performed based on the position of the target object on the first side of the vehicle. Causing the vehicle to travel within the first lane constraint and the second lane constraint such that the first distance is greater than the second distance may be revoked based on the input of the driver of the vehicle.

[0012]

Consistent with another disclosed embodiment, a method for navigating a vehicle is provided.

The method may include: using at least one image capturing device to acquire a plurality of images of a region near the vehicle; determining a first lane constraint on a first side of the vehicle based on the plurality of images; determining based on the plurality of images A second lane constraint on a second side of the vehicle opposite to the first side of the vehicle, wherein the first lane constraint and the second lane constraint define the lane in which the vehicle travels, and wherein the first lane constraint The distance corresponds to the distance between the first side of the vehicle and the first lane constraint, and the second distance corresponds to the distance between the second side of the vehicle and the second lane constraint Based on the plurality of images to determine whether the lane offset condition exists on the first side of the vehicle; if the lane offset condition exists on the first side of the vehicle, cause the vehicle Travel within the first lane constraint and the second lane constraint so that the first distance is greater than the second distance; based on the plurality of images, it is determined whether a lane offset condition exists in the first lane of the vehicle And if the lane offset condition exists on the second side of the vehicle, causing the vehicle to travel within the first lane constraint and the second lane constraint, such that the first distance is less than The second distance.

[0013]

Consistent with the disclosed embodiments, a driver assistance navigation system for a vehicle is provided.

The system may include: at least one image capturing device configured to acquire a plurality of images of an area near the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine the current travel lane among the plurality of available travel lanes based on the plurality of images; and if the current travel lane is consistent with a predetermined If the default lanes of travel are not the same, the vehicle will change lanes.

[0014]

Consistent with the disclosed embodiment, the predetermined default travel lane may be the rightmost lane among the plurality of available travel lanes.

The at least one processing device may be configured to activate a turn signal before causing the vehicle to change lanes. The at least one processing device may be configured to generate a voice notification before causing the vehicle to change lanes. The at least one processing device may be configured to generate a voice instruction before causing the vehicle to change lanes. The predetermined default lane of travel may be determined based on the plurality of images of the area near the vehicle, based on data received from a global positioning system, and/or based on input received from a user.

[0015]

Consistent with another disclosed embodiment, a vehicle may include: a main body; at least one image capturing device configured to acquire a plurality of images of an area near the vehicle; a data interface; and at least one processing Device.

The at least one processing device may be configured to: receive the plurality of images via the data interface; determine the current travel lane among the plurality of available travel lanes based on the plurality of images; and if the current travel lane is consistent with a predetermined If the default lanes of travel are not the same, the vehicle will change lanes.

[0016]

Consistent with the disclosed embodiment, the predetermined default travel lane may be the rightmost lane among the plurality of available travel lanes.

The at least one processing device may be configured to activate a turn signal before causing the vehicle to change lanes. The at least one processing device may be configured to generate a voice notification before causing the vehicle to change lanes. The at least one processing device may be configured to generate a voice instruction before causing the vehicle to change lanes. The predetermined default lane of travel may be determined based on the plurality of images of the area near the vehicle, based on data received from a global positioning system, and/or based on input received from a user.

[0017]

Consistent with another disclosed embodiment, a method for navigating a vehicle is provided.

The method may include: using at least one image capturing device to acquire a plurality of images of the area near the vehicle; determining the current travel lane among a plurality of available travel lanes based on the plurality of images; and if the current travel lane is consistent with a predetermined If the default lanes of travel are not the same, the vehicle will change lanes.

[0018]

Consistent with the disclosed embodiment, the predetermined default travel lane may be the rightmost lane among the plurality of available travel lanes.

The turn signal may be activated before causing the vehicle to change lanes. The predetermined default lane of travel may be determined based on the plurality of images of the area near the vehicle.

[0019]

Consistent with the disclosed embodiments, a driver assistance navigation system for a vehicle is provided.

The system may include: at least one image capturing device configured to acquire a plurality of images of an area near the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive

the plurality of images via the data interface; identify the curve to be navigated based on map data and vehicle position information; and at least based on the curve reflected in the map data A characteristic to determine the initial target speed of the vehicle; adjust the speed of the vehicle to the initial target speed; determine one or more observed characteristics of a curve based on the multiple images; The observed one or more characteristics of the road are used to determine the updated target speed; and the speed of the vehicle is adjusted to the updated target speed.

[0020]

Consistent with the disclosed embodiments, at least one characteristic of the curve may include an expected radius of curvature of the curve determined according to the map data.

The initial target speed may be determined based on the expected radius of curvature and the lateral acceleration limit of the curve. The lateral acceleration limit may be between about 0.2 m/s2 and 3.0 m/s2. Alternatively, the lateral acceleration limit may be between about 1.5 m/s2 and 3.0 m/s2. The lateral acceleration limit may be adjusted based on the road width determined by the plurality of images. The observed one or more characteristics of the curve may include the actual radius of curvature of the curve determined by analyzing the plurality of images. The updated target speed may be determined based on the actual radius of curvature and the lateral acceleration limit of the curve. The observed one or more characteristics of the curve may include a rate of change of the actual radius of curvature of the curve determined by analyzing the plurality of images. The observed one or more characteristics of the curve may include a curve inclination associated with the curve determined by analyzing the plurality of images. The observed one or more characteristics of the curve may include information related to the curve conveyed by one or more road signs determined by analyzing the plurality of images. The initial target speed may also be based on the focal length range of the at least one image capture device. The initial target speed may also be based on determining whether a traffic light is positioned along the curve. The initial target speed may also be based on determining whether the traffic light is within a predetermined distance behind the curve.

[0021]

Consistent with another disclosed embodiment, a vehicle may include: a main body; at least one image capturing device configured to acquire a plurality of images of an area near the vehicle; a data interface; and at least one processing Device.

The at least one processing device may be configured to: receive the plurality of images via the data interface; identify the curve to be navigated based on map data and vehicle position information; and at least based on the curve reflected in the map data A characteristic to determine the initial target speed of the vehicle; adjust the speed of the vehicle to the initial target speed; determine one or more observed characteristics of a curve based on the multiple images; The observed one or more characteristics of the road are used to determine the updated target speed; and the speed of the vehicle is adjusted to the updated target speed.

[0022]

Consistent with the disclosed embodiments, at least one characteristic of the curve may include an expected radius of curvature of the curve determined according to the map data.

The initial target speed may be determined based on the expected radius of curvature and the lateral acceleration limit of the curve. The lateral acceleration limit may be adjusted based on the road width determined by the plurality of images. The observed one or more characteristics of the curve may include the actual radius of curvature of the curve determined by analyzing the plurality of images. The updated target speed may be determined based on the actual radius of curvature and the lateral acceleration limit of the curve. The observed one or more characteristics of the curve may include a rate of change of the actual radius of curvature of the curve determined by analyzing the plurality of images. The observed one or more characteristics of the curve may include a curve inclination associated with the curve determined by analyzing the plurality of images. The observed one or more characteristics of the curve may include information related to the curve conveyed by one or more road signs determined by analyzing the plurality of images. The initial target speed may also be based on the focal length range of the at least one image capture device. The initial target speed may also be based on determining whether a traffic light is positioned along the curve. The initial target speed may also be based on determining whether the traffic light is within a predetermined distance behind the curve.

[0023]

Consistent with another disclosed embodiment, a method for navigating a vehicle is provided.

The method may include: using at least one image capturing device to acquire a plurality of images of the area near the vehicle; identifying the curve to be navigated based on map data and vehicle position information; and at least based on the curve reflected in the map data A characteristic to determine the initial target speed of the vehicle; adjust the speed of the vehicle to the initial target speed; determine one or more observed characteristics of a curve based on the multiple images; The observed one or more characteristics of the road are used to determine the updated target speed; and the speed of the vehicle is adjusted to the updated target speed.

[0024]

Consistent with the disclosed embodiments, a driver-assisted navigation system for a host vehicle is provided.

The system may include: at least one image capturing device configured to acquire a plurality of images of a region near the host vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine the first lane constraint on the first side of the host vehicle according to at least some of the plurality of images; The second lane constraint on the second side of the host vehicle opposite to the first side of the host vehicle is determined according to at least some of the images, wherein the first lane constraint and the first side A two-lane constraint defines the lane in which the host vehicle travels; causes the host vehicle to travel within the first lane constraint and the second lane constraint; locates the vehicle ahead in the multiple images; and based on the multiple images Determining at least one action taken by the preceding vehicle; and causing the host vehicle to imitate the at least one action of the preceding vehicle.

[0025]

Consistent with the disclosed embodiment, at least one of the first lane restriction and the second lane restriction may include a line marked on the road surface.

At least one of the first lane constraint and the second lane constraint may include a curb of a road surface. At least one of the first lane constraint and the second lane constraint may include a determined midpoint of the road surface width. The at least one action taken by the preceding vehicle may include a turn, and the at least one processing device may cause the host vehicle to perform a similar turn that starts at a location where the preceding vehicle initiates its turn. The at least one action taken by the preceding vehicle may include acceleration. The at least one action taken by the preceding vehicle may include braking. The at least one action taken by the preceding vehicle may include providing a first offset distance on the side of the preceding vehicle adjacent to the first lane constraint, the first offset distance being different from the preceding vehicle The second offset distance of the side adjacent to the second lane constraint. The at least one processing device may be further configured to determine whether the vehicle in front is turning at the intersection based on the plurality of images or based on map data and location information, and reject if it is determined that the vehicle in front is turning at the intersection imitating the turning of the preceding vehicle. The at least one processing device may be configured to determine whether the preceding vehicle is turning at the intersection based on the plurality of images or based on map data and location information, determining whether the preceding vehicle is changing lanes within the intersection, and if it is determined If the vehicle in front turns at an intersection and does not change lanes in the intersection, then imitate the turn of the vehicle in front, and if it is determined that the vehicle in front turns at the intersection but changes lanes within the intersection, then Refuse to imitate the turn of the preceding vehicle. The at least one processing device may be configured to acquire a first vehicle ahead based on the multiple images, acquire a second vehicle ahead based on the multiple images, and if the first vehicle ahead is determined based on the multiple images The turning radius of the path on which the vehicle continues to travel is smaller than the path followed by the second preceding vehicle, which causes the host vehicle to imitate the first preceding vehicle. The at least one processing device may be configured to determine the position information of the preceding vehicle and compare the position information with predetermined map data, by tracking the movement of the preceding vehicle relative to the predetermined map data as the preceding vehicle A navigation history is created, and based on the navigation history of the preceding vehicle, it is determined whether to cause the host vehicle to imitate the at least one action of the preceding vehicle. The at least one processing device may be configured to: when the navigation history indicates that the preceding vehicle has successfully navigated at least one forward turn, even when a turn is not detected based on the plurality of images, cause the host vehicle to imitate the turning of the vehicle in front.

[0026]

Consistent with another disclosed embodiment, a host vehicle may include: a main body; at least one image capturing device configured to acquire a plurality of images of an area near the host vehicle; a data interface; and at least A processing device.

The at least one processing device may be configured to: receive the plurality of images via the data interface; determine the first lane constraint on the first side of the host vehicle according to at least some of the plurality of images Determine the second lane constraint on the second side of the host vehicle opposite to the first side of the host vehicle according to at least some of the images, wherein the first lane constraint And the second lane constraint define the lane in which the host vehicle travels; cause the host vehicle to travel within the first lane constraint and the second lane constraint; locate the preceding vehicle in the multiple images; based on the multiple image to determine at least one action taken by the preceding vehicle; and causing the host vehicle to imitate the at least one action of the preceding vehicle.

[0027]

Consistent with the disclosed embodiment, at least one of the first lane restriction and the second lane restriction may include a line marked on the road surface.

At least one of the first lane constraint and the second lane constraint may include a curb of a road surface. At least one of the first lane constraint and the second lane constraint may include a determined midpoint of the road surface width. The at least one action taken by the preceding vehicle may include a turn, and the at least one processing device may cause the host vehicle to perform a similar turn that starts at a location where the preceding vehicle initiates its turn. The at least one action taken by the preceding vehicle may include acceleration. The at least one action taken by the preceding vehicle may include braking. The at least one action taken by the preceding vehicle may include providing a first offset distance on the side of the preceding vehicle adjacent to the first lane constraint, the first offset distance being different from the preceding vehicle The second offset distance of the side adjacent to the second lane constraint. The at least one processing device may be further configured to determine whether the vehicle in front is turning at the intersection based on the plurality of images or based on map data and location information, and reject if it is determined that the vehicle in front is turning at the intersection Imitating the turning of the preceding vehicle. The at least one processing device may be configured to determine whether the preceding vehicle is turning at the intersection based on the plurality of images or based on map data and location information, determining whether the preceding vehicle is changing lanes within the intersection, and if it is determined If the vehicle in front turns at an intersection and does not change lanes in the intersection, then imitate the turn of the vehicle in front, and if it is determined that the vehicle in front turns at the intersection but changes lanes within the intersection, then Refuse to imitate the turn of the preceding vehicle. The at least one processing device may be configured to acquire a first vehicle ahead based on the multiple images, acquire a second vehicle ahead based on the multiple images, and if the first vehicle ahead is determined based on the multiple images The turning radius of the path on which the vehicle continues to travel is smaller than the path followed by the second preceding vehicle, which causes the host vehicle to imitate the first preceding vehicle. The at least one processing device may be configured to determine the position information of the preceding vehicle and compare the position information with predetermined map data, by tracking the movement of the preceding vehicle relative to the predetermined map data as the preceding vehicle A navigation history is created, and based on the navigation history of the preceding vehicle, it is determined whether to cause the host vehicle to imitate the at least one action of the preceding vehicle. The at least one processing device may be configured to: when the navigation history indicates that the preceding vehicle has successfully navigated at least one forward turn, even when a turn is not detected based on the plurality of images, cause the host vehicle to imitate the turning of the vehicle in front.

[0028]

Consistent with another disclosed embodiment, a method for navigating a host vehicle is provided.

The method may include: using at least one image capturing device to acquire a plurality of images of the vicinity of the host vehicle; and determining the first lane on the first side of the host vehicle based on at least some of the images Constraints; determine a second lane constraint on a second side of the host vehicle opposite to the first side of the host vehicle according to at least some of the images, wherein the first lane The constraint and the second lane constraint define the lane in which the host vehicle travels; cause the host vehicle to travel within the first lane constraint and the second lane constraint; locate the preceding vehicle in the multiple images; based on the multiple images to determine at least one action taken by the preceding vehicle; and causing the host vehicle to imitate the at least one action of the preceding vehicle.

[0029]

Consistent with the disclosed embodiments, a driver-assisted navigation system for a user's vehicle is provided.

The system may include: at least one image capturing device configured to acquire a plurality of images of an area near the user's vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine the first lane constraint on the first side of the user's vehicle according to the plurality of images; The image determines the second lane constraint on the second side of the user vehicle opposite to the first side of the user vehicle, wherein the first lane constraint and the second lane constraint define the travel of the user vehicle Lane; acquiring a target vehicle based on the pluraility of images; if it is determined that the lane in which the target vehicle travels is different from the lane in which the user vehicle travels, the user vehicle is enabled to pass the target vehicle; based on the plurality of images To monitor the position of the target vehicle; and if it is determined that the target vehicle is entering the lane where the user vehicle is traveling, causing the user vehicle to abort the overtaking before completing the overtaking.

[0030]

Consistent with the disclosed embodiments, the step of causing the user's vehicle to suspend the overtaking may include causing the brake in the user's vehicle to be applied.

The step of causing the user vehicle to abort the overtaking may include generating a voice notification. The step of causing the user vehicle to stop the overtaking may include causing the user vehicle to stop accelerating. The step of monitoring the position of the target vehicle may include estimating the speed of the target vehicle. At least one of the first lane constraint and the second lane constraint may include a line marked on the road surface, a curb of the road surface, and/or a determined midpoint of the road surface width.

[0031]

Consistent with another disclosed embodiment, a user vehicle may include: at least one image capture device configured to acquire a plurality of images of a region near the user vehicle; a data interface; and at least one processing Device.

The at least one processing device may be configured to: receive the plurality of images via the data interface; determine the first lane constraint on the first side of the user's vehicle according to the plurality of images; The image determines a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first lane constraint and the second lane constraint define the travel of the user vehicle Lane; acquiring a target vehicle based on the plurality of images; if it is determined that the lane in which the target vehicle travels is different from the lane in which the user vehicle travels, the user vehicle is enabled to pass the target vehicle; based on the pluraility of images To monitor the position of the target vehicle; and if it is determined that the target vehicle is entering the lane where the user vehicle is traveling, causing the user vehicle to abort the overtaking before completing the overtaking.

[0032]

Consistent with the disclosed embodiments, causing the user's vehicle to abort the overtaking may include causing application of a brake in the user's vehicle.

Causing the user vehicle to abort the overtaking may include generating a voice notification. Causing the user vehicle to abort the overtaking may include causing the user vehicle to stop accelerating. Monitoring the location of the target vehicle may include estimating the speed of the target vehicle. At least one of the first lane constraint and the second lane constraint may include a line marked on the road surface, a curb of the road surface, and/or a determined midpoint of the road surface width.

[0033]

Consistent with another disclosed embodiment, a method for navigating a user's vehicle is provided.

The method may include: using at least one image capturing device to acquire a plurality of images of an area near the user's vehicle; determining a first lane constraint on the first side of the user's vehicle based on the plurality of images; Images determine a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first lane constraint and the second lane constraint define the user vehicle to travel Based on the multiple images to obtain the target vehicle; if it is determined that the lane in which the target vehicle travels is different from the lane in which the user vehicle travels, then the user vehicle can overtake the target vehicle; based on the multiple images to monitor the location of the

Tesla EX1002 Page 1175 IPR2025-00943

target vehicle; and if it is determined that the target vehicle is entering the lane where the user vehicle is traveling, causing the user vehicle to abort the overtaking before completing the overtaking.

[0034]

Consistent with the disclosed embodiments, the step of causing the user's vehicle to suspend the overtaking may include causing the brake in the user's vehicle to be applied.

The step of causing the user vehicle to stop the overtaking may include causing the user vehicle to stop accelerating. At least one of the first lane constraint and the second lane constraint may include a line marked on the road surface, a curb of the road surface, and/or a determined midpoint of the road surface width.

[0035]

Consistent with the disclosed embodiments, a driver-assisted navigation system for a user's vehicle is provided.

The system may include: at least one image capturing device configured to acquire a plurality of images of an area near the user's vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine the first lane constraint on the first side of the user's vehicle according to the plurality of images; The image determines the second lane constraint on the second side of the user vehicle opposite to the first side of the user vehicle, wherein the first lane constraint and the second lane constraint define the travel of the user vehicle Lane; determining whether an intruding vehicle approaches from one side of the user vehicle based on the plurality of images; and causing the user vehicle to maintain the current speed and travel within the first lane constraint and the second lane constraint, so that The first distance on the side of the user vehicle where the intruding vehicle is approaching is greater than the second distance.

[0036]

Consistent with the disclosed embodiment, the at least one processing device may be further configured to cause the user vehicle to take an evasive action if the analysis of the multiple images indicates that the intruding vehicle has crossed at least one collision threshold.

The avoidance action may include changing lanes, changing routes, and/or braking. The at least one collision threshold may include a minimum predetermined distance maintained between the user vehicle and the intruding vehicle. The at least one collision threshold may include the time to collision. The at least one processing device may be configured to cause the user vehicle to return to the center of the lane where the user vehicle originally traveled after completing the avoidance action. The at least one image capturing device may have a field of view of at least 140 degrees. The at least one processing device may be configured to cause the user's vehicle to return to a preset speed after braking in response to the intruding vehicle. The at least one processing device may be configured to cause the user vehicle to return to the lane location within the first lane constraint and the second lane constraint, so that after changing the route in response to the intruding vehicle, the user vehicle and the The first distance between first lane constraints is substantially equal to the second distance between the user vehicle and the second lane constraints.

[0037]

Consistent with another disclosed embodiment, a user vehicle may include: at least one image capture device configured to acquire a plurality of images of a region near the user vehicle; a data interface; and at least one processing Device.

The at least one processing device may be configured to: receive the plurality of images via the data interface; determine the first lane constraint on the first side of the user's vehicle according to the plurality of images; The image determines the second lane constraint on the second side of the user vehicle opposite to the first side of the user vehicle, wherein the first lane constraint and the second lane constraint define the travel of the user vehicle Lane; determining whether an intruding vehicle approaches from one side of the user vehicle based on the plurality of images; and causing the user vehicle to maintain the current speed and travel within the first lane constraint and the second lane constraint, so that The first distance on the side of the user vehicle where the intruding vehicle is approaching is greater than the second distance.

[0038]

IPR2025-00943 Tesla EX1002 Page 1177

Consistent with the disclosed embodiments, the at least one processing device may be further configured to cause the user vehicle to take an evasive action if the analysis of the multiple images indicates that the intruding vehicle has crossed at least one collision threshold.

The avoidance action may include changing lanes, changing routes, and/or braking. The at least one collision threshold may include a minimum predetermined distance maintained between the user vehicle and the intruding vehicle. The at least one collision threshold may include the time to collision. The at least one processing device may be configured to cause the user vehicle to return to the center of the lane where the user vehicle originally traveled after completing the avoidance action. The at least one image capturing device may have a field of view of at least 140 degrees. The at least one processing device may be configured to cause the user's vehicle to return to a preset speed after braking in response to the intruding vehicle. The at least one processing device may be configured to cause the user vehicle to return to the lane location within the first lane constraint and the second lane constraint, so that after changing the route in response to the intruding vehicle, the user vehicle and the The first distance between first lane constraints is substantially equal to the second distance between the user vehicle and the second lane constraints.

[0039]

Consistent with another disclosed embodiment, a method for navigating a user's vehicle is provided.

The method may include: using at least one image capturing device to acquire a plurality of images of an area near the user's vehicle; determining a first lane constraint on the first side of the user's vehicle based on the plurality of images; Images determine a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first lane constraint and the second lane constraint define the user vehicle to travel Determining whether an intruding vehicle approaches from one side of the user vehicle based on the multiple images; and causing the user vehicle to maintain the current speed and travel within the first lane constraint and the second lane constraint, The first distance on the side of the user's vehicle where the invading vehicle is approaching is made greater than the second distance.

[0040]

Consistent with the disclosed embodiments, a driver-assisted navigation system for a host vehicle is provided.

The system may include: at least one image capturing device configured to acquire a plurality of images of a region near the host vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; identify the target object in the plurality of images; monitor the movement of the target object and the host via the plurality of images The distance between the vehicle and the target object; an indicator that determines the interception time between the host vehicle and the target object based on the monitored movement and the distance between the host vehicle and the target object; and The response of the host vehicle is generated based on the comparison of the interception time with a plurality of predetermined interception thresholds.

[0041]

Consistent with the disclosed embodiment, the target object may be a target vehicle traveling in a lane different from the host vehicle.

The interception time may be based on an indicator of the speed of the tracked object determined by the plurality of images. The interception time may be based on an indicator of the acceleration of the tracked object determined by the plurality of images. The at least one processing device may be further configured to determine whether the target object is traveling in the same lane as the host vehicle, and if it is determined that the target object is traveling in a different lane from the host vehicle, abandon the The response of the host vehicle. Determining whether the target object is traveling in the same lane as the host vehicle may be based on the plurality of images. Determining whether the target object is traveling in the same lane as the host vehicle may be based on at least one of predetermined map data or yaw rate of the host vehicle. If the indicator of interception time indicates that the target object is outside the first interception threshold relative to the vehicle, the response may be based on an average motion profile associated with the target object. If the indicator of the interception time indicates that the target object is between the first interception threshold and the second interception threshold relative to the vehicle, the response may be based on the actual movement of the target object. If the indicator of interception time indicates that the target object is within the second interception threshold, the response may include an emergency avoidance action. The emergency avoidance action may include lane change and/or emergency braking.

[0042]

Consistent with another disclosed embodiment, a host vehicle may include: a main body; at least one image capturing device configured to acquire a plurality of images of an area near the host vehicle; a data interface; and at least A processing device.

The at least one processing device may be configured to: receive the plurality of images via the data interface; identify the target object in the plurality of images; monitor the movement of the target object and the host via the plurality of images The distance between the vehicle and the target object; an indicator that determines the interception time between the host vehicle and the target object based on the monitored movement and the distance between the host vehicle and the target object; and The response of the host vehicle is generated based on the comparison of the interception time with a plurality of predetermined interception thresholds.

[0043]

Consistent with the disclosed embodiment, the target object may be a target vehicle traveling in a lane different from the host vehicle.

The interception time may be based on an indicator of the speed of the tracked object determined by the plurality of images. The interception time may be based on an indicator of the acceleration of the tracked object determined by the plurality of images. The at least one processing device may be further configured to determine whether the target object is traveling in the same lane as the host vehicle, and if it is determined that the target object is traveling in a different lane from the host vehicle, abandon the The response of the host vehicle. Determining whether the target object is traveling in the same lane as the host vehicle may be based on the plurality of images. Determining whether the target object is traveling in the same lane as the host vehicle may be based on at least one of predetermined map data or yaw rate of the host vehicle. If the indicator of interception time indicates that the target object is outside the first interception threshold relative to the vehicle, the response may be based on an average motion profile associated with the target object. If the indicator of the interception time indicates that the target object is between the first interception threshold and the second interception threshold relative to the vehicle, the response may be based on the actual movement of the target object. If the indicator of interception time indicates that the target object is within the second interception threshold, the response may include an emergency avoidance action. The emergency avoidance action may include lane change and/or emergency braking.

[0044]

Consistent with another disclosed embodiment, a method for navigating a host vehicle may include: acquiring a plurality of images of an area near the host vehicle via at least one image capturing device; identifying a target object in the plurality of images; based on The plurality of images are used to monitor the movement of the target object and the distance between the host vehicle and the target object; determine the movement based on the monitored movement and the distance between the host vehicle and the target object An indicator of the interception time between the host vehicle and the target object; and generating the response of the host vehicle based on a comparison of the interception time with a plurality of predetermined interception thresholds.

[0045]

Consistent with the disclosed embodiments, a driver assistance navigation system for a vehicle is provided.

The system may include: at least one image capturing device configured to acquire a plurality of images of an area near the vehicle; a data interface; and at least one processing device.

The at least one processing device may be configured to: receive the plurality of images via the data interface; identify a trigger for stopping the vehicle based on analysis of the plurality of images; and the trigger based on the identification , Causing the vehicle to stop according to the braking profile, as determined based on the analysis of the plurality of images, the braking profile including a first section associated with a first deceleration rate, having a The second section of the second deceleration rate of the speed and the third section where the braking level is reduced when approaching the target stopping point.

[0046]

Consistent with the disclosed embodiments, reaching the target stopping point may be accompanied by the release of brake pressure.

The trigger may include an object. The object may include another vehicle in front of the vehicle. The trigger may include a traffic light signal, a stop sign, and/or a speed limit sign. The at least one processing device may be further configured to provide a voice notification identifying the trigger. The at least one processing device may be configured to generate a voice notification indicating that the vehicle is braking.

[0047]

Consistent with another disclosed embodiment, a vehicle may include: a main body; at least one image capturing device configured to acquire a plurality of images of an area near the vehicle; a data interface; and at least one processing Device.

The at least one processing device may be configured to: receive the plurality of images via the data interface; identify a trigger for stopping the vehicle based on analysis of the plurality of images; and the trigger based on the identification A device to cause the vehicle to stop according to a braking profile, as determined based on the analysis of the plurality of images, the braking profile includes a first section associated with a first deceleration rate, and has a The second section of the second deceleration rate of the speed and the third section where the braking level is reduced when approaching the target stopping point.

[0048]

Consistent with the disclosed embodiments, reaching the target stopping point may be accompanied by the release of brake pressure.

The trigger may include an object. The object may include another vehicle in front of the vehicle. The trigger may include a traffic light signal, a stop sign, and/or a speed limit sign. The at least one processing device may be further configured to provide a voice notification identifying the trigger. The at least one processing device may be configured to generate a voice notification indicating that the vehicle is braking.

[0049]

Consistent with another disclosed embodiment, a method for navigating a vehicle is provided.

The method may include: acquiring a plurality of images of an area near the vehicle via at least one image capturing device; identifying a trigger for stopping the vehicle based on analysis of the plurality of images; and the trigger based on the identification , Causing the vehicle to stop according to the braking profile, as determined based on the analysis of the plurality of images, the braking profile includes a first section associated with a first deceleration rate, and has a first section that is less than the first deceleration rate The second section of the second deceleration rate and the third section that reduces the braking level when approaching the target stopping point.

[0050]

Consistent with other disclosed embodiments, a non-transitory computer-readable storage medium may store program instructions, which are executed by at least one processing device and execute any method described herein.

[0051]

The foregoing general description and the following detailed description are only exemplary and illustrative and do not limit the claims.

[0052]

Description of the drawings

[0053]

The drawings incorporated in and forming a part of the present disclosure illustrate various disclosed embodiments.

In the picture:

[0054]

Figure 1 is a schematic representation of an exemplary system consistent with the disclosed embodiments.

[0055]

Figure 2A is a schematic side view representation of an exemplary vehicle including a system consistent with the disclosed embodiments.

[0056]

Figure 2B is a schematic top view representation of the vehicle and system shown in Figure 2A consistent with the disclosed embodiment.

[0057]

Figure 2C is a schematic top view representation of another embodiment of a vehicle including a system consistent with the disclosed embodiment.

[0058]

Figure 2D is a schematic top view representation of another embodiment of a vehicle including a system consistent with the disclosed embodiment.

[0059]

Figure 2E is a schematic top view representation of yet another embodiment of a vehicle including a system consistent with the disclosed embodiment.

[0060]

Figure 2F is a schematic representation of an exemplary vehicle control system consistent with the disclosed embodiments.

[0061]

Figure 3A is a schematic representation of the interior of a vehicle consistent with the disclosed embodiment, including a rearview mirror and a user interface for the vehicle imaging system.

[0062]

3B is an illustration of an example of a camera mount configured to be positioned behind the rearview mirror and against the windshield of the vehicle consistent with the disclosed embodiment.

[0063]

FIG. 3C is an illustration of a camera mount with a different viewing angle shown in FIG. 3B consistent with the disclosed embodiment.

[0064]

3D is an illustration of an example of a camera mount configured to be positioned behind the rearview mirror and against the windshield of the vehicle consistent with the disclosed embodiment.

[0065]

4 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

[0066]

FIG. 5A is a flowchart showing an exemplary process for generating one or more navigation responses based on monocular image analysis consistent with the disclosed embodiments.

[0067]

FIG. 5B is a flowchart showing an exemplary process for detecting one or more vehicles and/or pedestrians in a set of images consistent with the disclosed embodiments.

[0068]

FIG. 5C is a flowchart showing an exemplary process for detecting road signs and/or lane geometric information in a set of images consistent with the disclosed embodiments.

[0069]

FIG. 5D is a flowchart showing an exemplary process for detecting traffic lights in a set of images consistent with the disclosed embodiments.

[0070]

FIG. 5E is a flowchart showing an exemplary process for generating one or more navigation responses based on a vehicle path consistent with the disclosed embodiment.

[0071]

FIG. 5F is a flowchart showing an exemplary process for determining whether the vehicle in front changes lanes, consistent with the disclosed embodiment.

[0072]

FIG. 6 is a flowchart showing an exemplary process for generating one or more navigation responses based on stereoscopic image analysis consistent with the disclosed embodiments.

[0073]

FIG. 7 is a flowchart showing an exemplary process for generating one or more navigation responses based on three sets of image analysis consistent with the disclosed embodiments.

[0074]

Figure 8A is a schematic representation of an exemplary vehicle traveling within a first lane constraint and a second lane constraint consistent with the disclosed embodiment, where the lane offset condition exists on the first vehicle side.

[0075]

FIG. 8B is a schematic representation of an exemplary vehicle traveling within the first lane constraint and the second lane constraint consistent with the disclosed embodiment, where the lane offset condition exists on the second vehicle side.

[0076]

FIG. 8C is a schematic representation of an exemplary vehicle traveling within a first lane constraint and a second lane constraint on a curved road consistent with the disclosed embodiment.

[0077]

8D is a schematic representation of an exemplary vehicle traveling within the first lane constraint and the second lane constraint consistent with the disclosed embodiment, wherein the lane offset condition exists on both the first vehicle side and the second vehicle side .

[0078]

Figure 9 is a schematic representation of the memory of an exemplary navigation system consistent with the disclosed embodiments.

[0079]

FIG. 10 is a flowchart of an exemplary method for navigating a vehicle consistent with the disclosed embodiment.

[0080]

Figure 11 is a schematic representation of an exemplary vehicle including a navigation system traveling in a default lane consistent with the disclosed embodiments.

[0081]

Figure 12A is a schematic representation of the memory of an exemplary navigation system consistent with the disclosed embodiments.

[0082]

FIG. 12B is a flowchart of an exemplary process for navigating a vehicle consistent with the disclosed embodiment.

[0083]

Figures 13A and 13B are schematic representations of an exemplary vehicle approaching and navigating a curve having one or more characteristics consistent with the disclosed embodiments.

[0084]

FIG. 14 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

[0085]

FIG. 15 is a flowchart showing an exemplary process for controlling the vehicle speed based on the detected curve and the observed curve characteristic consistent with the disclosed embodiment.

[0086]

16A to 16D are schematic representations of a host vehicle that imitates one or more actions of the preceding vehicle consistent with the disclosed embodiment.

[0087]

FIG. 17 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

[0088]

FIG. 18 is a flowchart showing an exemplary process for causing the host vehicle to imitate one or more actions of the preceding vehicle consistent with the disclosed embodiment.

[0089]

FIG. 19 is a flowchart showing an exemplary process for causing the host vehicle to refuse to imitate the turn of the preceding vehicle consistent with the disclosed embodiment.

[0090]

FIG. 20 is a flowchart showing an exemplary process for causing the host vehicle to imitate or refuse the turn of the preceding vehicle consistent with the disclosed embodiment.

[0091]

FIG. 21 is a flowchart showing another exemplary process for causing the host vehicle to imitate one or more actions of the preceding vehicle consistent with the disclosed embodiment.

[0092]

FIG. 22 is a flowchart showing an exemplary process for causing the host vehicle to imitate one or more actions of the preceding vehicle based on the navigation history, consistent with the disclosed embodiment.

[0093]

FIG. 23 is a schematic representation of an exemplary vehicle traveling within a first lane constraint and a second lane constraint on a road along which another vehicle is traveling, consistent with the disclosed embodiment.

[0094]

FIG. 24A is a schematic representation of a memory that stores instructions for navigating the vehicle in an invading vehicle consistent with the disclosed embodiment.

[0095]

FIG. 24B is a flowchart of an exemplary process for navigating a vehicle among an invading vehicle consistent with the disclosed embodiment.

[0096]

FIG. 25 is a schematic representation of an exemplary vehicle traveling within a first lane constraint and a second lane constraint on a road along which another vehicle is traveling, consistent with the disclosed embodiment.

[0097]

Fig. 26 is a schematic representation of a memory storing instructions for detecting and responding to lateral intrusion of vehicle traffic consistent with the disclosed embodiment.

[0098]

FIG. 27 is a flowchart of an exemplary process for detecting and responding to lateral intrusion of vehicle traffic consistent with the disclosed embodiments.

[0099]

28A and 28B are schematic representations of generating a host vehicle response consistent with the disclosed embodiment.

[0100]

FIG. 29 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

[0101]

FIG. 30 is a flowchart showing an exemplary process for generating a host vehicle response consistent with the disclosed embodiment.

[0102]

FIG. 31 is a flowchart showing an exemplary process for refusing to generate a host vehicle response consistent with the disclosed embodiment.

[0103]

FIG. 32A is a schematic representation of an exemplary vehicle in which an object exists in the vicinity of the vehicle consistent with the disclosed embodiment.

[0104]

Figure 32B illustrates an exemplary braking profile consistent with the disclosed embodiment.

[0105]

Figure 33 is a schematic representation of the memory of an exemplary navigation system consistent with the disclosed embodiments.

[0106]

FIG. 34 is a flowchart of an exemplary method for implementing a braking profile while navigating a vehicle, consistent with the disclosed embodiment.

[0107]

Detailed ways

[0108]

The following detailed description refers to the accompanying drawings.

Where possible, the same reference signs are used in the figures and in the following description to refer to the same or similar parts.

Although several illustrative embodiments have been described herein, modifications, adaptations, and other implementations are possible.

For example, substitutions, additions, or modifications can be made to the components illustrated in the figures, and the illustrative methods described herein can be modified by substituting, reordering, removing, or adding steps to the disclosed methods.

Therefore, the following detailed description is not limited to the disclosed embodiments and examples.

Instead, the proper scope is defined by the appended claims.

[0109]

The disclosed embodiments provide systems and methods that use cameras to provide autonomous navigation features.

In various embodiments, the system may include one, two, or more cameras that monitor the environment of the vehicle.

In one embodiment, the system can monitor and adjust the free space between the vehicle and the lane boundary in which the vehicle travels.

In another embodiment, the system can select a specific lane as the default lane used by the vehicle while traveling.

In another embodiment, the system can control the speed of the vehicle in different scenarios, such as when turning.

In yet another embodiment, the system can mimic the actions of the preceding vehicle.

In yet another embodiment, the system may monitor the target vehicle and enable the vehicle to pass the target vehicle under certain conditions (for example, if the lane in which the target vehicle travels is different from the lane in which the vehicle travels).

In yet another embodiment, the system may generate a natural feel in response to a lateral intrusion into a vehicle, such as a vehicle attempting to enter the lane in which the vehicle is traveling.

[0110]

Figure 1 is a block diagram representation of a system 100 consistent with an exemplary disclosed embodiment.

Depending on the requirements of a particular implementation, the system 100 may include various components.

In some embodiments, the system 100 may include a processing unit 110, an image acquisition unit 120, a position sensor 130, one or more memory units 140, 150, a map database 160, and a user interface 170.

The processing unit 110 may include one or more processing devices.

In some embodiments, the processing unit 110 may include an application processor 180, an image processor 190, or any other suitable processing device.

Similarly, depending on the requirements of a particular application, the image acquisition unit 120 may include any number of image acquisition devices and components.

In some embodiments, the image capturing unit 120 may include one or more image capturing devices (for example, cameras), such as the image capturing device 122, the image capturing device 124, and the image capturing device 126.

The system 100 may further include a data interface 128 that communicatively connects the processing device 110 and the image acquisition device 120.

For example, the data interface 128 may include any one or more wired and/or wireless links for sending the image data acquired by the image acquisition device 120 to the processing unit 110.

[0111]

Both the application processor 180 and the image processor 190 may include various types of processing devices.

For example, either or both of the application processor 180 and the image processor 190 may include a microprocessor, a pre-processor (such as an image pre-processor), a graphics processor, a central processing unit (CPU), support circuits, Digital signal processor, integrated circuit, memory or any other type of device suitable for running application programs and image processing and analysis.

In some embodiments, the application processor 180 and/or the image processor 190 may include any type of single-core or multi-core processors, mobile device microcontrollers, central processing units, and so on.

Various processing devices may be used, including, for example, processors available from manufacturers such as , and may include various structures (for example, x86 processors, , etc.).

[0112]

In some embodiments, the application processor 180 and/or the image processor 190 may include any EyeQ series processor chips available from .

These processor designs each include multiple processing units with local memory and instruction sets.

Such a processor may include a video input for receiving image data from a plurality of image sensors, and may also include a video output function.

In one example, uses 90 nanometer micron technology operating at 332Mhz.

The structure of is as follows: two floating-point, hyper-threaded 32-bit RISC CPU (core), five visual computing engines (VCE), three vector microcode processors Denali 64-bit mobile DDR controller, 128-bit internal Supersonic interconnection, dual 16-bit video input controller and 18-bit video output controller, 16-channel DMA and several peripherals.

The MIPS34K CPU manages five VCEs, three VMPTMs and DMA, a second MIPS34K CPU and multi-channel DMA, and other peripherals.

Five VCEs, three and MIPS34K CPUs can perform the intensive visual calculations required by multifunctional bundled applications.

In another example, , which is a third-generation processor and six times more powerful than , can be used in the disclosed implementation.

[0113]

Although Figure 1 depicts two separate processing devices included in the processing unit 110, more or fewer processing devices may be used.

For example, in some embodiments, a single processing device may be used to complete the tasks of the application processor 180 and the image processor 190.

In other embodiments, these tasks may be performed by more than two processing devices.

[0114]

The processing unit 110 may include various types of devices.

For example, the processing unit 110 may include various devices, such as a controller, an image preprocessor, a central processing unit (CPU), a support circuit, a digital signal processor, an integrated circuit, a memory, or any other type for image processing and analysis. installation.

The image preprocessor may include a video processor for capturing, digitizing, and processing images from the image sensor.

The CPU may include any number of microcontrollers or microprocessors.

The support circuit can be any number of circuits generally known in the art, including cache, power supply, clock, and input and output circuits.

The memory can store software that, when executed by the processor, controls the operation of the system.

The memory may include a database and image processing software.

Tesla EX1002 Page 1186 IPR2025-00943

The memory may include any number of random access memory, read-only memory, flash memory, disk drives, optical storage media, tape storage media, removable storage media, and other types of storage media.

In one example, the memory may be detached from the processing unit 110.

In another example, the memory may be integrated into the processing unit 110.

[0115]

Each memory 140, 150 may include software instructions that, when executed by a processor (for example, the application processor 180 and/or the image processor 190), can control various aspects of the operation of the system 100.

These memory units may include various databases and image processing software.

The memory unit may include random access memory, read-only memory, flash memory, disk drives, optical storage media, tape storage media, removable storage media, and/or any other types of storage media.

In some embodiments, the memory units 140, 150 may be separated from the application processor 180 and/or the image processor 190.

In other embodiments, these memory units may be integrated into the application processor 180 and/or the image processor 190.

[0116]

The position sensor 130 may include any type of device suitable for determining a location associated with at least one component of the system 100.

In some embodiments, the location sensor 130 may include a GPS receiver.

Such a receiver can determine the user's location and speed by processing signals propagated by GPS satellites.

The position information from the position sensor 130 can be used by the application processor 180 and/or the image processor 190.

[0117]

The user interface 170 may include any device suitable for providing information to or for receiving input from one or more users of the system 100.

In some embodiments, the user interface 170 may include user input devices, including, for example, a touch screen, a microphone, a keyboard, a pointing device, a tracking wheel, a camera, a knob, a button, and so on.

In the case of such an input device, the user can provide information input or commands to the system 100 through the following means: typing instructions or information, providing voice commands, using buttons, pointers, or eye-tracking functions to select menu options on the screen or using Any other suitable technology for communicating information with the system 100.

[0118]

The user interface 170 may be equipped with one or more processing devices configured to provide information to or receive information from the user and process the information for use by the application processor 180, for example.

In some embodiments, such a processing device can execute instructions to recognize and track eye movements, receive and interpret voice commands, recognize and interpret touches and/or gestures made on the touch screen, and respond to keyboard input or menu selection ,and many more.

In some embodiments, the user interface 170 may include a display, a speaker, a haptic device, and/or any other device for providing output information to the user.

[0119]

The map database 160 may include any type of database for storing map data useful to the system 100.

Tesla EX1002 Page 1187 IPR2025-00943

In some embodiments, the map database 160 may include location-related data in a reference coordinate system for various items, including roads, water features, geographic features, businesses, points of interest, restaurants, gas stations, etc. .

The map database 160 can store not only the locations of such items, but also descriptive words related to these items (including, for example, names associated with any stored features).

In some embodiments, the map database 160 may be physically located with other components of the system 100.

Alternatively or in addition, the map database 160 or a portion thereof may be remotely located relative to other components of the system 100 (e.g., the processing unit 110).

In such an embodiment, the information from the map database 160 may be downloaded through a wired or wireless data connection of the network (for example, through a cellular network and/or the Internet, etc.).

[0120]

The image capture devices 122, 124, and 126 may each include any type of device suitable for capturing at least one image from the environment.

Moreover, any number of image capture devices can be used to capture images for input to the image processor.

Some embodiments may include only a single image capture device, while other embodiments may include two, three, or even more than four image capture devices.

The image capturing devices 122, 124, and 126 will be further described below with reference to FIGS. 2B to 2E.

[0121]

The system 100 or its various components can be integrated into a variety of different platforms.

In some embodiments, the system 100 may be included on a vehicle 200 as shown in FIG. 2A.

For example, as described above with respect to FIG. 1, the vehicle 200 may be equipped with the processing unit 110 and any other components of the system 100.

While in some embodiments, the vehicle 200 may only be equipped with a single image capture device (eg, a camera), in other embodiments such as those discussed in conjunction with FIGS. 2B to 2E, multiple image capture devices may be used.

For example, any one of the image capturing devices 122 and 124 of the vehicle 200 as shown in FIG. 2A may be part of an ADAS (Advanced Driver Assistance System) imaging group.

[0122]

The image capturing device included on the vehicle 200 as a part of the image capturing unit 120 may be located at any suitable location.

In some embodiments, as shown in FIGS. 2A to 2E and 3A to 3C, the image capture device 122 may be located near the rearview mirror.

This location can provide a line of sight (similar to the line of sight of the driver of the vehicle 200), which can help determine what is visible to the driver and what is invisible.

The image capture device 122 can be positioned anywhere near the rear view mirror, but placing the image capture device 122 on the side of the mirror that is on the driver can further help obtain an image representing the driver's field of view and/or line of sight.

[0123]

For the image capturing device of the image capturing unit 120, other locations may also be used.

For example, the image capture device 124 may be located on or in the bumper of the vehicle 200.

Such a location may be particularly suitable for image capture devices with a wide field of view.

The line of sight of the image capturing device located on the bumper may be different from the line of sight of the driver, and therefore, the bumper image capturing device and the driver may not always see the same object.

Image capture devices (e.g., image capture devices 122, 124, and 126) may also be located in other locations.

For example, the image capture device can be located on or in one or two side mirrors of the vehicle 200, on the roof of the vehicle 200, on the hood of the vehicle 200, on the trunk of the vehicle 200, or on the side of the vehicle 200. On any window of the vehicle 200, positioned behind or in front of any window of the vehicle 200, installed in or near the light map on the front and/or back of the vehicle 200, and so on.

[0124]

In addition to the image capture device, the vehicle 200 may include various other components of the system 100.

For example, the processing unit 110 may be included in the vehicle 200, and may be integrated or separated from an engine control unit (ECU) of the vehicle.

The vehicle 200 may also be equipped with a position sensor 130 such as a GPS receiver, and may also include a map database 160 and memory units 140 and 150.

[0125]

Figure 2A is a schematic side view representation of an exemplary vehicle imaging system consistent with the disclosed embodiments.

Fig. 2B is a schematic top view illustration of the embodiment shown in Fig. 2A. As illustrated in FIG. 2B, the disclosed embodiment may include a vehicle 200 that includes a system 100 in its body, the system 100 having a first image capture positioned near the rearview mirror and/or near the driver of the vehicle 200 The device 122, the second image capturing device 124 positioned on or in the bumper area (for example, one of the bumper areas 210) of the vehicle 200, and the processing unit 110.

[0126]

As illustrated in FIG. 2C, both the image capture devices 122 and 124 may be positioned near the rearview mirror and/or near the driver of the vehicle 200.

In addition, although two image capture devices 122 and 124 are shown in FIGS. 2B and 2C, it should be understood that other embodiments may include more than two image capture devices. For example, in the embodiments shown in FIGS. 2D and 2E, the first image capturing device 122, the second image capturing device 124, and the third image capturing device 126 are included in the system 100 of the vehicle 200.

[0127]

As illustrated in FIG. 2D, the image capture device 122 may be located near the rearview mirror and/or near the driver of the vehicle 200, and the image capture devices 124 and 126 may be located in the bumper area of the vehicle 200 (e.g., bumper One of the regions 210) above or in the middle.

And as shown in FIG. 2E, the image capturing devices 122, 124, and 126 may be positioned near the rearview mirror and/or near the driver's seat of the vehicle 200. The disclosed embodiments are not limited to any specific number and configuration of image capture devices, and the image capture devices may be located in any suitable location within and/or on the vehicle 200.

[0128]

It should be understood that the disclosed embodiments are not limited to vehicles and can be applied in other contexts.

It should also be understood that the disclosed embodiments are not limited to a particular type of vehicle 200 and can be applied to all types of vehicles, including automobiles, trucks, trailers, and other types of vehicles.

[0129]

The first image capture device 122 may include any suitable type of image capture device.

The image capture device 122 may include an optical axis. In one example, the image capture device 122 may include an Aptina M9V024WVGA sensor with a global shutter. In other embodiments, the image capture device 122 may provide a resolution of 1280×960 pixels and may include a rolling shutter. The image capturing device 122 may include various optical elements. In some embodiments, for example, one or more lenses may be included to provide a desired focal length and field of view for the image capture device. In some embodiments, the image capture device 122 may be associated with a 6mm lens or a 12mm lens. In some embodiments, as illustrated in FIG. 2D, the image capture device 122 may be configured to capture an image having a desired field of view (FOV) 202. For example, the image capturing device 122 may be configured to have a conventional FOV, such as in the range of 40 degrees to 56 degrees, including 46 degrees FOV, 50 degrees FOV, and 52 degrees FOV or more. Alternatively, the image capturing device 122 may be configured to have a narrow FOV in the range of 23 degrees to 40 degrees, such as a 28-degree FOV or a 36-degree FOV. In addition, the image capturing device 122 may be configured to have a wide FOV in the range of 100 degrees to 180 degrees. In some embodiments, the image capture device 122 may include a wide-angle bumper camera or have a FOV up to 180 degrees.

[0130]

The first image capturing device 122 may acquire a plurality of first images with respect to a scene associated with the vehicle 200.

Each of the plurality of first images can be acquired as a series of image scan lines that can be captured using a rolling shutter. Each scan line may include multiple pixels.

[0131]

The first image capture device 122 may have a scan rate associated with acquiring each of the scan lines of the first series of images.

The scan rate may refer to a rate at which the image sensor can acquire image data associated with each pixel included in a specific scan line.

[0132]

For example, the image capture devices 122, 124, and 126 may include any suitable type and number of image sensors, including CCD sensors or CMOS sensors.

In one embodiment, a CMOS image sensor can be used together with a rolling shutter, so that each pixel in a row is read one at a time, and the scanning of the row is performed on a line by line basis until the entire image frame has been captured . In some embodiments, the row can be captured sequentially from top to bottom with respect to the frame.

[0133]

Using a rolling shutter can cause pixels in different rows to be exposed and captured at different times, which can cause skew and other image artifacts in the captured image frame.

On the other hand, when the image capture device 122 is configured to operate with a global or synchronized shutter, all pixels can be exposed for the same amount of time and in a common exposure period. As a result, the image data of the frames collected from the system using the global shutter represents a snapshot of the entire FOV (such as FOV 202) at a specific time. In contrast, in a rolling shutter application, each line in the frame is exposed and the data is captured at different times. As a result, the moving object may appear to be distorted in an image capture device with a rolling shutter. This phenomenon will be described in more detail below.

[0134]

The second image capturing device 124 and the third image capturing device 126 may be any types of image capturing devices.

Like the first image capture device 122, each of the image capture devices 124 and 126 may include an optical axis. In one embodiment, each of the image capture devices 124 and 126 may include an Aptina M9V024 WVGA sensor with a global shutter. Alternatively, each of the image capture devices 124 and 126 may include a rolling shutter. Like the image capturing

device 122, the image capturing devices 124 and 126 may be configured to include various lenses and optical elements. In some embodiments, the lenses associated with image capture devices 124 and 126 may provide FOVs (such as FOVs 204 and 206) that are equal to or narrower than the FOV associated with image capture device 122 (such as FOV 202). For example, the image capture devices 124 and 126 may have a FOV of 40 degrees, 30 degrees, 26 degrees, 23 degrees, 20 degrees or less.

[0135]

The image capturing devices 124 and 126 may acquire a plurality of second images and third images with respect to the scene associated with the vehicle 200.

Each of the plurality of second images and third images can be acquired as scan lines of the second series and the third series that can be captured using a rolling shutter. Each scan line or row can have multiple pixels. The image capturing devices 124 and 126 may have a second scan rate and a third scan rate associated with acquiring each image scan line included in the second series and the third series.

[0136]

Each image capture device 122, 124, and 126 can be positioned in any suitable position and orientation relative to the vehicle 200.

The relative positioning of the image capture devices 122, 124, and 126 can be selected to help merge the information acquired from the image capture devices. For example, in some embodiments, the FOV associated with the image capture device 124 (such as FOV 204) may be partially related to the FOV associated with the image capture device 122 (such as FOV 202) and the FOV associated with the image capture device 126 (such as FOV 206). Or completely overlap.

[0137]

The image capture devices 122, 124, and 126 may be located at any suitable relative height on the vehicle 200.

In one example, there may be a height difference between the image capture devices 122, 124, and 126, which may provide sufficient disparity information to enable stereo analysis. For example, as shown in FIG. 2A, the two image capturing devices 122 and 124 are at different heights. For example, there may also be a lateral displacement difference between the image capturing devices 122, 124, and 126, which provides additional disparity information for the stereo analysis of the processing unit 110. The lateral displacement difference can be recorded as dx, as shown in Figure 2C and Figure 2D. In some embodiments, front-to-back displacement (eg, range displacement) may exist between the image capture devices 122, 124, and 126. For example, the image capture device 122 may be located 0.5 meters to more than 2 meters behind the image capture device 124 and/or the image capture device 126. This type of displacement enables one of the image capture devices to cover the potential blind spot of the other image capture device.

[0138]

The image capture device 122 may have any suitable resolution function (for example, the number of pixels associated with the image sensor), and the resolution of the image sensor associated with the image capture device 122 may be greater than, less than, or equal to that of the image capture devices 124 and 126 The resolution of the associated image sensor.

In some embodiments, the image sensors associated with image capture device 122 and/or image capture devices 124 and 126 may have a resolution of 640×480, 1024×768, 1280×960, or any other suitable resolution.

[0139]

The frame rate (for example, the rate at which the image capture device acquires a set of pixel data for an image frame before moving to capture the pixel data associated with the next image frame) may be controllable.

The frame rate associated with image capture device 122 may be greater than, less than, or equal to the frame rate associated with image capture devices 124 and 126. The frame rate associated with the image capture devices 122, 124, and 126 may depend on various factors that may affect the timing of the frame rate. For example, one or more of the image capture devices 122, 124, and 126 may include information imposed before or after acquiring image data associated with one or more pixels of the image sensor in the image capture device 122, 124, and/or 126. Optional pixel delay period. Generally, the image data

IPR2025-00943
Tesla EX1002 Page 1191

corresponding to each pixel can be acquired according to the clock rate of the device (for example, one pixel per clock cycle). In addition, in embodiments that include rolling shutters, one or more of the image capture devices 122, 124, and 126 may be included in acquiring a row of pixels associated with the image sensor in the image capture devices 122, 124, and/or 126. An optional horizontal blanking period imposed before or after the image data. Further, one or more of the image capture devices 122, 124, and/or 126 may include an optional vertical blanking period that is imposed before or after the image data associated with the image frames of the image capture devices 122, 124, and 126 are acquired.

[0140]

These timing controls can synchronize the frame rates associated with image capture devices 122, 124, and 126, even when the line scan rate of each image capture device is different.

In addition, as will be discussed in more detail below, among other factors (e.g., image sensor resolution, maximum line scan rate, etc.), these optional timing controls can make the FOV of the image capture device 122 and image capture The image capture of the area where one or more FOVs of the devices 124 and 126 overlap is synchronized, even when the field of view of the image capture device 122 is different from the FOV of the image capture devices 124 and 126.

[0141]

The frame rate timing of the image capture devices 122, 124, and 126 may depend on the resolution of the associated image sensor.

For example, assuming that two devices have similar line scan rates, if one device includes an image sensor with a resolution of 640x480 and the other device includes an image sensor with a resolution of 1280x960, it will take more time to obtain High-resolution sensor image data frame.

[0142]

Another factor that can affect the timing of image data acquisition by the image capture devices 122, 124, and 126 is the maximum line scan rate.

For example, acquiring one line of image data of the image sensors included in the image capture devices 122, 124, and 126 will require some minimum amount of time. Assuming no pixel delay period is added, this minimum amount of time to acquire one line of image data will be related to the maximum line scan rate of a particular device. Devices that provide a higher maximum line scan rate potentially provide a larger frame rate than devices that have a smaller maximum line scan rate. In some embodiments, the maximum line scan rate of one or more of the image capture devices 124 and 126 may be greater than the maximum line scan rate associated with the image capture device 122. In some embodiments, the maximum line scan rate of the image capture device 124 and/or 126 may be 1.25 times, 1.5 times, 1.75 times, or more than 2 times the maximum line scan rate of the image capture device 122.

[0143]

In another embodiment, the image capture devices 122, 124, and 126 may have the same maximum line scan rate, but the image capture device 122 may operate at a scan rate less than or equal to its maximum scan rate.

The system may be configured such that one or more of the image capture devices 124 and 126 operate at a line scan rate equal to the line scan rate of the image capture device 122. In other examples, the system may be configured such that the line scan rate of the image capture device 124 and/or the image capture device 126 may be 1.25 times, 1.5 times, 1.75 times, or more than 2 times the line scan rate of the image capture device 122 .

[0144]

In some embodiments, the image capture devices 122, 124, and 126 may be asymmetrical.

That is, they may include cameras with different fields of view (FOV) and focal lengths. For example, the field of view of the image capture devices 122, 124, and 126 may include any desired area relative to the environment of the vehicle 200. In some embodiments, one or more of the image capture devices 122, 124, and 126 may be configured to acquire image data from the environment in front of the vehicle 200, behind the vehicle 200, from the side of the vehicle 200, or a combination thereof.

[0145]

Further, the focal length associated with each image capture device 122, 124, and/or 126 may be optional (for example, by including an appropriate lens, etc.) so that each device captures a desired distance range relative to the vehicle 200 The image of the object at the location.

For example, in some embodiments, the image capture devices 122, 124, and 126 may acquire images of close-up objects within a few meters of the vehicle. The image capturing devices 122, 124, and 126 may also be configured to acquire images of objects farther away from the vehicle (for example, 25m, 50m, 100m, 150m or more). Further, the focal lengths of the image capturing devices 122, 124, and 126 can be selected so that one image capturing device (for example, the image capturing device 122) can acquire images of objects relatively close to the vehicle (for example, within 10 m or within 20 m), and other images Capture devices (eg, image capture devices 124 and 126) can capture images of objects farther from the vehicle 200 (eg, greater than 20m, 50m, 100m, 150m, etc.).

[0146]

According to some embodiments, the FOV of one or more image capture devices 122, 124, and 126 may have a wide angle.

For example, especially for image capturing devices 122, 124, and 126 that can be used to capture images of areas near the vehicle 200, it may be advantageous to have a FOV of 140 degrees. For example, the image capture device 122 may be used to capture an image of an area on the right or left of the vehicle 200, and in such an embodiment, the image capture device 122 may desirably have a wide FOV (e.g., at least 140 degrees).

[0147]

The field of view associated with each of the image capture devices 122, 124, and 126 may depend on the corresponding focal length.

For example, as the focal length increases, the corresponding field of view decreases.

[0148]

The image capture devices 122, 124, and 126 may be configured to have any suitable field of view.

In a specific example, the image capture device 122 may have a horizontal FOV of 46 degrees, the image capture device 124 may have a horizontal FOV of 23 degrees, and the image capture device 126 may have a horizontal FOV between 23 degrees and 46 degrees. In another example, the image capture device 122 may have a horizontal FOV of 52 degrees, the image capture device 124 may have a horizontal FOV of 26 degrees, and the image capture device 126 may have a horizontal FOV between 26 degrees and 52 degrees. In some embodiments, the ratio of the FOV of the image capture device 122 to the FOV of the image capture device 124 and/or the image capture device 126 may vary from 1.5 to 2.0. In other embodiments, the ratio can vary between 1.25 and 2.25.

[0149]

The system 100 may be configured such that the field of view of the image capture device 122 and the field of view of the image capture device 124 and/or the image capture device 126 at least partially or completely overlap.

In some embodiments, the system 100 may be configured such that the field of view of the image capture devices 124 and 126, for example, falls within (eg, narrower than) the field of view of the image capture device 122 and shares a common field with the field of view of the image capture device 122. center. In other embodiments, the image capture devices 122, 124, and 126 may capture adjacent FOVs, or may have their FOVs partially overlap. In some embodiments, the field of view of the image capture devices 122, 124, and 126 may be aligned so that the center of the narrower FOV image capture device 124 and/or 126 may be located in the lower half of the field of view of the wider FOV device 122 .

[0150]

Figure 2F is a schematic representation of an exemplary vehicle control system consistent with the disclosed embodiments.

As indicated in FIG. 2F, the vehicle 200 may include a throttle system 220, a braking system 230, and a steering system 240. The system 100 can be connected to one of the throttle system 220, the braking system 230, and the steering system 240 via one or more data links (for example, any wired and/or wireless one or more links used to transmit data). Or more provide inputs (e.g., control signals). For example, based on the analysis of the images acquired by the image capture devices 122, 124, and/or 126, the system 100 may provide control signals to one or more of the throttle system 220, the braking system 230, and the steering system 240 to (For example, by causing acceleration, turning, lane shifting, etc.) to navigate the vehicle 200. Further, the system 100 may receive input from one or more of the throttle system 220, the braking system 230, and the steering system 240 to indicate the operating conditions of the vehicle 200 (for example, speed, whether the vehicle 200 is braking and/or turning, and many more). Below, further details are provided in conjunction with FIGS. 4-7.

[0151]

As shown in FIG. 3A, the vehicle 200 may also include a user interface 170 for interacting with the driver or passenger of the vehicle 200.

For example, in a vehicle application, the user interface 170 may include a touch screen 320, a knob 330, a button 340, and a microphone 350. The driver or passenger of the vehicle 200 may also use a handle (for example, located on or near the steering column of the vehicle 200, including, for example, a turn signal handle), buttons (for example, located on the steering wheel of the vehicle 200), etc., to interact with the system 100 . In some embodiments, the microphone 350 may be positioned adjacent to the rear view mirror 310. Similarly, in some embodiments, the image capture device 122 may be located near the rear view mirror 310. In some embodiments, the user interface 170 may also include one or more speakers 360 (eg, speakers of a vehicle audio system). For example, the system 100 may provide various notifications (e.g., prompts) via the speaker 360.

[0152]

3B to 3D are illustrations of an exemplary camera mount 370 consistent with the disclosed embodiments. The camera mount 370 is configured to be positioned behind the rearview mirror (for example, the rearview mirror 310) and against the vehicle block Wind glass.

As shown in FIG. 3B, the camera mount 370 may include image capture devices 122, 124, and 126. The image capturing devices 124 and 126 may be positioned behind the light shield 380, which may be flush with the windshield of the vehicle and include a combination of films and/or anti-reflective materials. For example, the shading plate 380 may be positioned so that it is aligned against a vehicle windshield with a matching slope. In some embodiments, for example, as depicted in FIG. 3D, each of the image capture devices 122, 124, and 126 may be positioned behind the light shield 380. The disclosed embodiments are not limited to the image capturing devices 122, 124, and 126, the camera mount 370, and the light shielding plate 380 of any specific configuration. FIG. 3C is an illustration of the camera mount 370 shown in FIG. 3B from a front view.

[0153]

As those skilled in the art who benefit from the present disclosure will understand, many variations and/or modifications can be made to the previously disclosed embodiments.

For example, not all components are critical to the operation of the system 100. Further, any component can be located in any suitable part of the system 100, and the components can be rearranged into various configurations while providing the functionality of the disclosed embodiments. Therefore, the foregoing configuration is an example and regardless of the configuration discussed above, the system 100 can provide a wide range of functionality for analyzing the surrounding environment of the vehicle 200 and navigating the vehicle 200 in response to the analysis.

[0154]

As discussed in further detail below and consistent with various disclosed embodiments, the system 100 may provide various features related to autonomous driving and/or driver assistance technology.

For example, the system 100 may analyze image data, location data (for example, GPS location information), map data, speed data, and/or data from sensors included in the vehicle 200. The system 100 may collect data from, for example, the image acquisition unit 120, the position sensor 130, and other sensors for analysis. Further, the system 100 may analyze the collected data to determine whether the vehicle 200 should take certain actions, and then automatically take certain actions without human intervention. For example, when the vehicle 200 is navigating without human intervention, the system 100 can

IPR2025-00943 Tesla EX1002 Page 1195

automatically control the vehicle 200 (for example, by sending a control signal to one or more of the throttle system 220, the braking system 230, and the steering system 240). Braking, acceleration and/or steering. Further, the system 100 may analyze the collected data and issue a warning, and/or prompt the vehicle driver and occupant based on the analysis of the collected data.

[0155]

Further, consistent with the disclosed embodiments, the functionality provided by the system 100 may cause the vehicle 200 to take different actions to navigate the vehicle 200 within the lane and/or relative to other vehicles and/or objects.

For example, the system 100 may adjust the positioning of the vehicle 200 relative to the lane in which the vehicle 200 travels and/or relative to objects located near the vehicle 200, select a specific lane for the vehicle 200 to use while traveling, and respond to intrusive vehicles such as The vehicle attempting to enter the lane in which the vehicle 200 is traveling) takes action. In addition, the system 100 can control the speed of the vehicle 200 in different scenarios, such as when the vehicle 200 makes a turn. The system 100 may further cause the vehicle 200 to imitate the actions of the preceding vehicle or monitor the target vehicle, and guide the vehicle 200 to overtake the target vehicle. Additional details regarding the various implementations provided by the system 100 are provided below.

[0156]

Forward multi-imaging system

[0157]

As discussed above, the system 100 may provide driving assistance functionality using a multi-camera system.

The multi-camera system may use one or more cameras facing the forward direction of the vehicle.

In other embodiments, the multi-camera system may include one or more cameras facing the side of the vehicle or the rear of the vehicle. In one embodiment, for example, the system 100 may use a dual camera imaging system, in which the first camera and the second camera (for example, the image capture devices 122 and 124) may be positioned in front of and/or in front of the vehicle (for example, the vehicle 200). side. The field of view of the first camera may be larger, smaller than, or partially overlap with the field of view of the second camera. In addition, the first camera can be connected to the first image processor to perform monocular image analysis of the image provided by the first camera, and the second camera can be connected to the second image processor to perform the image provided by the second camera Monocular image analysis. The outputs (for example, processed information) of the first image processor and the second image processor may be combined. In some embodiments, the second image processor may receive images from both the first camera and the second camera to perform stereo analysis. In another embodiment, the system 100 may use a three-camera imaging system, where each camera has a different field of view. Therefore, such a system can make decisions based on information obtained from objects located at varying distances in both the front and the sides of the vehicle. Reference to monocular image analysis may refer to an instance in which image analysis is performed based on images captured from a single observation point (for example, from a single camera). Stereoscopic image analysis may refer to an instance in which image analysis is performed based on two or more images captured with one or more variations of image capture parameters. For example, a captured image suitable for performing stereoscopic image analysis may include an image captured using two or more different positions, from different fields of view, using different focal lengths, combining parallax information, and so on.

[0158]

For example, in one embodiment, the system 100 may implement a three-camera configuration using image capture devices 122-126.

In such a configuration, the image capture device 122 can provide a narrow field of view (for example, 34 degrees, or other values selected from the range of about 20 degrees to 45 degrees, etc.), and the image capture device 124 can provide a wide field of view (for example, , 150 degrees or other values selected from the range of about 100 degrees to about 180 degrees), and the image capture device 126 may provide an intermediate field of view (for example, 46 degrees or other values selected from the range of about 35 degrees to about 60 degrees). In some embodiments, the image capture device 126 may act as a main camera or a main camera. The image capture devices 122-126 may be positioned behind the rear view mirror 310 and positioned substantially side by side (e.g , 6 cm spacing). Further, in some embodiments, as discussed above, one or more of

the image capture devices 122-126 may be installed behind a light shield 380 that is flush with the windshield of the vehicle 200. Such shading can minimize the impact of any reflections from the vehicle on the image capturing devices 122-126.

[0159]

In another embodiment, as discussed above in connection with FIGS. 3B and 3C, a wide field of view camera (e.g., the image capture device 124 in the above example) may be installed lower than a narrow main field of view camera (e.g., the above example The image devices 122 and 126 in the).

This configuration can provide free line of sight from a wide field of view camera. In order to reduce reflections, the camera may be installed close to the windshield of the vehicle 200, and may include a polarizer on the camera to suppress reflected light.

[0160]

The three-camera system can provide certain performance characteristics.

For example, some embodiments may include the ability to verify the detection of an object by one camera based on the detection result of another camera. In the three-camera configuration discussed above, the processing unit 110 may include, for example, three processing devices (e.g., three EyeQ series processor chips, as discussed above), each of which is dedicated to processing by the image capture device 122 -One or more of the captured images in 126.

[0161]

In the three-camera system, the first processing device can receive images from both the main camera and the narrow field of view camera, and perform the visual processing of the narrow FOV camera to detect other vehicles, pedestrians, lane markings, traffic signs, and traffic lights, for example And other road objects.

Further, the first processing device may calculate the pixel gap between the images of the main camera and the narrow camera and create a three-dimensional reconstruction of the environment of the vehicle 200. Then, the first processing device may combine the three-dimensional reconstruction with the three-dimensional map data, or with the three-dimensional information calculated based on the information of another camera.

[0162]

The second processing device may receive images from the main camera and perform visual processing to detect other vehicles, pedestrians, lane markings, traffic signs, traffic lights, and other road objects.

In addition, the second processing device may calculate the camera displacement, and based on the displacement, calculate the pixel gap between consecutive images and create a three-dimensional reconstruction of the scene (for example, according to the structure of the motion). The second processing device may send the structure from the motion-based three-dimensional reconstruction to the first processing device to be combined with the stereoscopic three-dimensional image.

[0163]

The third processing device may receive images from a wide FOV camera and process the images to detect vehicles, pedestrians, lane markings, traffic signs, traffic lights, and other road objects.

The third processing device may further execute additional processing instructions to analyze the image to identify objects moving in the image, such as vehicles that change lanes, pedestrians, and so on.

[0164]

In some embodiments, enabling image-based information streams to be independently captured and processed can provide opportunities for system redundancy.

For example, such redundancy may include using a first image capturing device and images processed from the device to verify and/or supplement information obtained by capturing and processing at least image information from the second image capturing device.

[0165]

In some embodiments, the system 100 may use two image capture devices (eg, image capture devices 122 and 124) to provide navigation assistance to the vehicle 200, and a third image capture device (eg, image capture device 126) to provide Redundant and verify the analysis of the data received from the other two image capture devices.

For example, in such a configuration, the image capture devices 122 and 124 can provide images for the stereo analysis of the system 100 to navigate the vehicle 200, and the image capture device 126 can provide images for the monocular analysis of the system 100 to provide redundancy and verify information obtained from images captured by the image capture device 122 and/or the image capture device 124. That is, the image capture device 126 (and corresponding processing device) may be considered to provide a redundant subsystem to provide a check on the analysis obtained from the image capture devices 122 and 124 (for example, provide an automatic emergency braking (AEB) system) .

[0166]

Those skilled in the art will recognize that the above camera configuration, camera placement, number of cameras, camera location, etc. are just examples.

These components and other components described with respect to the overall system can be assembled and used in a variety of different configurations without departing from the scope of the disclosed embodiments. Below are further details on the use of multi-camera systems to provide driver assistance and/or autonomous vehicle functionality.

[0167]

4 is an exemplary functional block diagram of the memory 140 and/or 150 consistent with the disclosed embodiments. The memory 140 and/or 150 may store/program instructions for performing one or more operations.

Although the following refers to the memory 140, those skilled in the art will recognize that the instructions can be stored in the memory 140 and/or 150.

[0168]

As shown in FIG. 4, the memory 140 can store a monocular image analysis module 402, a stereo image analysis module 404, a speed and acceleration module 406, and a navigation response module 408.

The disclosed embodiments are not limited to the memory 140 of any particular configuration. Further, the application processor 180 and/or the image processor 190 can execute instructions stored in any of the modules 402-408 included in the memory 140. Those skilled in the art will understand that references to the processing unit 110 in the following discussion may refer to the application processor 180 and the image processor 190 individually or collectively. Therefore, the steps of any of the following processes may be executed by one or more processing devices.

[0169]

In one embodiment, the monocular image analysis module 402 may store instructions (such as computer vision software). When executed by the processing unit 110, the instructions execute a set of images acquired by one of the image capture devices 122, 124, and 126. Monocular image analysis.

In some embodiments, the processing unit 110 may combine information from a set of images with additional sensory information (for example, information from radar) to perform monocular image analysis. As described below in conjunction with FIGS. 5A to 5D, the monocular image analysis module 402 may include a set of features for detecting a set of images (such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, Dangerous objects and any other characteristics associated with the vehicle environment) instructions. Based on this analysis, the system 100 (for example, via the processing unit 110) can generate one or more navigation responses in the vehicle 200, such as turning, lane shifting, acceleration changes, etc., which are discussed in conjunction with the navigation response module 408 below.

[0170]

In one embodiment, the stereoscopic image analysis module 404 may store instructions (such as computer vision software). When executed by the processing unit 110, the instructions are executed by a combination of image capture devices selected from any image capture device 122, 124, and 126. And the acquired stereo image analysis of the first group and the second group of images.

In some embodiments, the processing unit 110 may combine information from the first and second sets of images with additional sensory information (for example, information from radar) to perform stereoscopic image analysis. For example, the stereoscopic image analysis module 404 may include instructions for performing stereoscopic image analysis based on the first set of images acquired by the image capture device 124 and the second set of images acquired by the image capture device 126. As described below in conjunction with FIG. 6, the stereo image analysis module 404 may include a set of features for detecting the first and second sets of images (such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic Lights, dangerous objects, etc.) instructions. Based on this analysis, the processing unit 110 may generate one or more navigation responses in the vehicle 200, such as turning, lane shifting, acceleration changes, etc., which are discussed below in conjunction with the navigation response module 408.

[0171]

In one embodiment, the speed and acceleration module 406 may store software configured to analyze data received from one or more computer electrical devices of the vehicle 200 that are configured to cause the speed and/or acceleration of the vehicle 200 Change.

For example, the processing unit 110 may execute instructions associated with the speed and acceleration module 406 to calculate the target speed of the vehicle 200 based on data obtained from executing the monocular image analysis module 402 and/or the stereo image analysis module 404. For example, such data may include: target position, speed and/or acceleration; position and/or speed of vehicle 200 relative to nearby vehicles, pedestrians or road objects; position information of vehicle 200 relative to lane markings of the road, and so on. In addition, the processing unit 110 may calculate the vehicle based on sensory input (for example, information from radar) and input from other systems of the vehicle 200 (such as the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200). A target speed of 200. Based on the calculated target speed, the processing unit 110 may send electronic signals to the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200 to be triggered by, for example, physically depressing the brake or releasing the accelerator of the vehicle 200 Changes in speed and/or acceleration.

[0172]

In one embodiment, the navigation response module 408 may store software executable by the processing unit 110 to determine the desired navigation response based on data obtained from the execution of the monocular image analysis module 402 and/or the stereo image analysis module 404.

Such data may include location and speed information associated with nearby vehicles, pedestrians, and road objects, target location information of the vehicle 200, and the like. In addition, in some embodiments, the navigation response may be (partially or completely) based on map data, the predetermined location of the vehicle 200, and/or the execution of the vehicle 200 and the slave monocular image analysis module 402 and/or the stereo image analysis module 404. The relative velocity or relative acceleration between the detected one or more objects. The navigation response module 408 may also determine the desired navigation response based on sensory input (e.g., information from radar) and input from other systems of the vehicle 200 (such as the throttle system 220, braking system 230, and steering system 240 of the vehicle 200) . Based on the desired navigation response, the processing unit 110 may send electronic signals to the throttle system 220, the braking system 230, and the steering system 240 of the vehicle 200 to trigger the desired navigation response, for example, by turning the steering wheel of the vehicle 200 to achieve the predetermined Angle rotation. In some embodiments, the processing unit 110 may use the output of the navigation response module 408 (eg, desired navigation response) as an input to execute the speed and acceleration module 406 to calculate the speed change of the vehicle 200.

[0173]

FIG. 5A is a flowchart illustrating an exemplary process 500A for generating one or more navigation responses based on monocular image analysis consistent with the disclosed embodiments.

In step 510, the processing unit 110 may receive a plurality of images via the data interface 128 between the processing unit 110 and the image acquisition unit 120. For example, a camera included in the image acquisition unit 120 (such as an image

IPR2025-00943 Tesla EX1002 Page 1199

capture device 122 having a field of view 202) can capture multiple images of an area in front of the vehicle 200 (for example, or the side or back of the vehicle) and be connected via data. The image is sent to the processing unit 110 (for example, digital, wired, USB, wireless, Bluetooth, etc.). In step 520, the processing unit 110 may execute the monocular image analysis module 402 to analyze multiple images, which is described in further detail below in conjunction with FIGS. 5B to 5D. By performing the analysis, the processing unit 110 can detect a set of features in a set of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, and so on.

[0174]

In step 520, the processing unit 110 may also execute the monocular image analysis module 402 to detect various road hazards, for example, a part of a truck tire, a fallen road sign, loose goods, small animals, and so on.

Road hazards may have changes in structure, shape, size, and color, which may make detecting such hazards more challenging. In some embodiments, the processing unit 110 can execute the monocular image analysis module 402 to perform multi-frame analysis on multiple images to detect road hazards. For example, the processing unit 110 may estimate the camera motion between consecutive image frames, and calculate the pixel gap between the frames to construct a three-dimensional map of the road. Then, the processing unit 110 may use the three-dimensional map to detect the road surface and the hazards existing above the road surface.

[0175]

In step 530, the processing unit 110 may execute the navigation response module 408 to generate one or more navigation responses of the vehicle 200 based on the analysis performed in step 520 and the technique described above in conjunction with FIG. 4.

For example, the navigation response may include turns, lane shifts, acceleration changes, and so on. In some embodiments, the processing unit 110 may use data obtained from the execution speed and acceleration module 406 to generate one or more navigation responses. In addition, multiple navigation responses can occur simultaneously, sequentially, or in any combination thereof. For example, the processing unit 110 may cause the vehicle 200 to change a lane, and then accelerate by sequentially sending control signals to the steering system 240 and the throttle system 220 of the vehicle 200, for example. Alternatively, the processing unit 110 may cause the vehicle 200 to brake while simultaneously changing lanes by, for example, sending a control signal to the braking system 230 and the steering system 240 of the vehicle 200 at the same time.

[0176]

Figure 5B is a flowchart illustrating an exemplary process 500B for detecting one or more vehicles and/or pedestrians in a set of images consistent with the disclosed embodiments.

The processing unit 110 may execute the monocular image analysis module 402 to implement the process 500B. In step 540, the processing unit 110 may determine a set of candidate objects representing possible vehicles and/or pedestrians. For example, the processing unit 110 may scan one or more images, compare the images with one or more predetermined patterns, and identify the possibility that the object of interest (for example, a vehicle, a pedestrian, or a part thereof) may be included in each image. Place. The predetermined pattern can be designed in such a way as to achieve a high rate of "false hits" and a low rate of "misses". For example, the processing unit 110 may use the similarity with a predetermined pattern low threshold to identify the candidate object as a possible vehicle or pedestrian. Doing so may allow the processing unit 110 to reduce the probability of not hitting (e.g., not identifying) a candidate object representing a vehicle or pedestrian.

[0177]

In step 542, the processing unit 110 may filter this set of candidate objects to exclude certain candidates (for example, irrelevant or less relevant objects) based on the classification criteria.

Such criteria can be derived from various attributes associated with object types stored in a database (e.g., a database stored in memory 140). The attributes may include object shape, size, texture, position (for example, relative to vehicle 200), and the like. Thus, the processing unit 110 may use one or more sets of criteria to reject erroneous candidates from a set of candidate objects.

[0178]

In step 544, the processing unit 110 may analyze multiple frames of the image to determine whether an object in a set of candidate objects represents a vehicle and/or a pedestrian.

For example, the processing unit 110 may track the detected candidate object across consecutive frames and accumulate frame-by-frame data (eg, size, position relative to the vehicle 200, etc.) associated with the detected object. In addition, the processing unit 110 may estimate the parameters of the detected object and compare the frame-by-frame position data of the object with the predicted position.

[0179]

In step 546, the processing unit 110 may construct a set of measurement results of the detected object.

For example, such measurement results may include position, velocity, and acceleration values (relative to the vehicle 200) associated with the detected object. In some embodiments, the processing unit 110 may construct measurement results based on the following techniques: based on estimation techniques using a series of time-based observations (such as Kalman filter or linear quadratic estimation (LQE)); and/or based on different objects Modeling data available for types (for example, cars, trucks, pedestrians, bicycles, road signs, etc.). The Kalman filter may be based on a measurement result of the scale of the object, where the scale measurement is proportional to the time to collision (for example, the amount of time the vehicle 200 will reach the object). Thus, by performing steps 540-546, the processing unit 110 can identify vehicles and pedestrians appearing in a set of captured images and obtain information (for example, position, speed, size) associated with the vehicles and pedestrians. Based on the identified and obtained information, the processing unit 110 may generate one or more navigation responses of the vehicle 200, which are described above in conjunction with FIG. 5A.

[0180]

In step 548, the processing unit 110 may perform optical flow analysis of one or more images to reduce the probability of detecting "false hits" and not hitting candidate objects representing vehicles or pedestrians.

Optical flow analysis may refer to, for example, analyzing the motion pattern of the vehicle 200 in one or more images associated with other vehicles and pedestrians, and distinguishing it from road surface motion. The processing unit 110 may calculate the motion of the candidate object by observing different positions of the object across multiple image frames (captured at different times). The processing unit 110 may use the position and time values as inputs for calculating the mathematical model of the motion of the candidate object. Thus, optical flow analysis can provide another method for detecting vehicles and pedestrians near the vehicle 200. The processing unit 110 can be combined with steps 540-546 to perform optical flow analysis, provide redundancy for detecting vehicles and pedestrians, and increase the reliability of the system 100.

[0181]

5C is a flowchart illustrating an exemplary process 500C for detecting road signs and/or lane geometric information in a set of images consistent with the disclosed embodiments.

The processing unit 110 may execute the monocular image analysis module 402 to implement the process 500C. In step 550, the processing unit 110 may detect a group of objects by scanning one or more images. In order to detect a section of lane markings, lane geometry information, and other related road signs, the processing unit 110 may filter a set of objects to exclude these objects that are determined to be irrelevant (for example, slight potholes, small rocks, etc.). In step 552, the processing unit 110 may group the sections detected in step 550 that belong to the same road signs or lane markings. Based on the grouping, the processing unit 110 can build a model representing the detected section, such as a mathematical model.

[0182]

In step 554, the processing unit 110 may construct a set of measurement results associated with the detected section.

In some embodiments, the processing unit 110 may create a projection of the detected segment from the image plane to the real world plane. The projection can be characterized using a third polynomial with coefficients corresponding to physical properties such as position, slope, curvature, and curvature derived from the detected road. When generating the projection, the processing unit 110 may consider road surface changes and pitch and yaw rates associated with the vehicle 200. In addition, the processing unit 110 may model the road elevation by analyzing the position and motion cues existing on the road surface.

Further, the processing unit 110 may estimate the pitch and yaw rate associated with the vehicle 200 by tracking a set of feature points in one or more images.

[0183]

In step 556, the processing unit 110 may perform multi-frame analysis by, for example, tracking the detected section across consecutive image frames and accumulating frame-by-frame data associated with the detected section.

When the processing unit 110 performs multi-frame analysis, the set of measurement results constructed in step 554 may become more reliable and associated with a higher confidence level. Thus, by performing steps 550-556, the processing unit 110 can identify road signs appearing in a set of captured images and obtain lane geometry information. Based on the identification and the obtained information, the processing unit 110 can generate one or more navigation responses of the vehicle 200, which are described above in conjunction with FIG. 5A.

[0184]

In step 558, the processing unit 110 may consider additional information sources to further establish a safety model of the vehicle 200 in the context of its surrounding environment.

The processing unit 110 may use a safety model to define the following background, that is, the system 100 may adopt a safe manner to perform autonomous control of the vehicle 200. In order to establish a safety model, in some embodiments, the processing unit 110 may consider the position and movement of other vehicles, the curbs and obstacles of the detected road, and/or the overall road shape extracted from map data (such as the data of the map database 160) describe. By considering additional information sources, the processing unit 110 can provide redundancy to detect road signs and lane geometry and increase the reliability of the system 100.

[0185]

FIG. 5D is a flowchart showing an exemplary process 500D for detecting traffic lights in a set of images consistent with the disclosed embodiments.

The processing unit 110 may execute the monocular image analysis module 402 to implement the process 500D. In step 560, the processing unit 110 may scan a set of images and identify objects appearing at locations in the images that may contain traffic lights. For example, the processing unit 110 may filter the recognized objects to construct a set of candidate objects, and exclude those objects that are unlikely to correspond to traffic lights. Filtering can be done based on various attributes associated with traffic lights, such as shape, size, texture, position (e.g., relative to vehicle 200, etc.). Such attributes can be based on multiple examples of traffic lights and traffic control signals and stored in a database. In some embodiments, the processing unit 110 may perform multi-frame analysis on a set of candidate objects reflecting possible traffic lights. For example, the processing unit 110 may track candidate objects across consecutive image frames, estimate the real-world position of the candidate objects, and filter out those objects that are moving (which cannot be traffic lights). In some embodiments, the processing unit 110 may perform color analysis of the candidate object, and identify the relative position of the detected color that appears inside the possible traffic light.

[0186]

In step 562, the processing unit 110 may analyze the geometry of the hub.

The analysis can be based on any combination of: (i) the number of lanes detected on either side of the vehicle 200; (ii) markings detected on the road (such as arrow markings); and (iii) from map data (such as The data of the map database 160) extract the description of the hub. The processing unit 110 may use the information obtained from the execution of the monocular analysis module 402 for analysis. In addition, the processing unit 110 may determine the correspondence between the traffic light detected in step 560 and the lane appearing in the vicinity of the vehicle 200.

[0187]

When the vehicle 200 approaches the hub, in step 564, the processing unit 110 may update the confidence level associated with the analyzed hub geometry and the detected traffic light.

For example, compared with the number of traffic lights actually appearing at the hub, the estimated number of traffic lights appearing at the hub may affect the confidence level. Thus, based on the confidence level, the processing unit 110 may authorize the driver of the vehicle 200 to perform control in order to improve safety conditions. By performing steps 560-564, the processing unit 110 can identify traffic lights appearing in a set of captured images and analyze the hub geometry information. Based on the identification and analysis, the processing unit 110 can generate one or more navigation responses of the vehicle 200, which are described above in conjunction with FIG. 5A.

[0188]

5E is a flowchart illustrating an exemplary process 500E for generating one or more navigation responses in the vehicle 200 based on the vehicle path consistent with the disclosed embodiments.

In step 570, the processing unit 110 may construct an initial vehicle path associated with the vehicle 200. The vehicle path may be represented by a set of points expressed in coordinates (x, z), and the distance di between two points of the set of points may fall within the range of 1 meter to 5 meters. In one embodiment, the processing unit 110 may use two polynomials such as left and right road polynomials to construct the initial vehicle path. The processing unit 110 may calculate the geometric midpoint between the two polynomials, and if there is an offset, offset each point included in the resulting vehicle path by a predetermined offset (for example, smart lane offset), (Zero offset may correspond to traveling in the middle of the lane). The direction of the offset may be perpendicular to the section between any two points in the vehicle path. In another embodiment, the processing unit 110 may use a polynomial and the estimated lane width to offset each point of the vehicle path by: half of the estimated lane width plus a predetermined offset (for example, smart lane offset) .

[0189]

In step 572, the processing unit 110 may update the vehicle path constructed in step 570.

The processing unit 110 may use a higher resolution to reconstruct the vehicle path constructed in step 570 so that the distance dk between two points representing a set of points of the vehicle path is smaller than the distance di described above. For example, the distance dk may fall within the range of 0.1 meters to 0.3 meters. The processing unit 110 may use a parabolic spline algorithm to reconstruct the vehicle path, which may generate a cumulative distance vector S corresponding to the total length of the vehicle path (ie, based on a set of points representing the vehicle path).

[0190]

In step 574, the processing unit 110 may determine a look-ahead point (expressed in coordinates (xl, zl)) based on the updated vehicle path constructed in step 572.

The processing unit 110 may extract a look-ahead point from the accumulated distance vector S, and the look-ahead point may be associated with the look-ahead distance and the look-ahead time. The lower limit range can be from 10 meters to 20 meters, the forward distance can be calculated as the product of the speed of the vehicle 200 and the forward time. For example, as the speed of the vehicle 200 decreases, the look-ahead distance may also decrease (e.g., until the look-ahead distance reaches a lower limit). The look-ahead time, which may range from 0.5 seconds to 1.5 seconds, may be inversely proportional to the gain of one or more control loops associated with generating a navigation response in the vehicle 200 (such as tracking the heading error of the control loop). For example, the gain of the heading error of the tracking control loop may depend on the bandwidth of the yaw rate loop, the steering actuator loop, the lateral dynamics of the vehicle, and so on. Therefore, the larger the gain of the heading error of the tracking control loop, the shorter the look-ahead time.

[0191]

In step 576, the processing unit 110 may determine the heading error and the yaw rate command based on the look-ahead point determined in step 574.

The processing unit 110 may determine the heading error by calculating the arctangent of the look-ahead point (for example, arctan(xl/zl)). The processing unit 110 may determine the yaw rate command as the product of the heading error and the high-level control gain. If the look-ahead distance is not at the lower limit, the high-level control gain can be equal to: (2/look-ahead time). In other ways, the high-level control gain can be equal to: (2 times the speed of the vehicle 200/forward distance).

[0192]

FIG. 5F is a flowchart illustrating an exemplary process 500F for determining whether the vehicle in front changes lanes, consistent with the disclosed embodiment.

In step 580, the processing unit 110 may determine navigation information associated with the preceding vehicle (for example, a vehicle traveling before the vehicle 200). For example, the processing unit 110 may use the technique described above in conjunction with FIGS. 5A and 5B to determine the position, speed (for example, direction and speed), and/or acceleration of the preceding vehicle. The processing unit 110 may also use the technique described above in conjunction with FIG. 5E to determine one or more road polynomials, forward-looking points (associated with the vehicle 200), and/or slow-moving traces (for example, a set of points describing the path taken by the preceding vehicle).).

[0193]

In step 582, the processing unit 110 may analyze the navigation information determined in step 580.

In one embodiment, the processing unit 110 may calculate the distance between the slow-moving trail and the road polynomial (for example, along the trail). If the variance of the distance along the trail exceeds a predetermined threshold (e.g. 0.1 m to 0.2 m on a straight road, 0.3 m to 0.4 m on a moderately curved road, and 0.5 m to 0.6 m on a sharply curved road) , The processing unit 110 may determine that the vehicle in front may change lanes. In a case where it is detected that there are multiple vehicles traveling in front of the vehicle 200, the processing unit 110 may compare the slow-moving traces associated with each vehicle. Based on the comparison, the processing unit 110 may determine that a vehicle whose slow-moving trace does not match the slow-moving trace of other vehicles may change lanes. The processing unit 110 may additionally compare the curvature of the slow-moving trace (associated with the preceding vehicle) with the expected curvature of the road section traveled by the preceding vehicle. The expected curvature can be extracted from map data (for example, data of the map database 160), road polynomials, slow-moving traces of other vehicles, existing knowledge about the road, and the like. If the curvature difference between the slow moving trace and the expected curvature of the road section exceeds a predetermined threshold, the processing unit 110 may determine that the vehicle in front may change lanes.

[0194]

In another embodiment, the processing unit 110 may compare the instantaneous position of the preceding vehicle with the look-ahead point (associated with the vehicle 200) within a specific time period (for example, 0.5 seconds to 1.5 seconds).

If the distance between the instantaneous position of the preceding vehicle and the forward-looking point changes during a specific time period, and the cumulative sum of the changes exceeds a predetermined threshold (for example, 0.3 m to 0.4 m on a straight road, 0.7 m to 0.8 on a moderately curved road Meters, and 1.3 meters to 1.7 meters on a sharply curved road), the processing unit 110 may determine that the vehicle in front may change lanes. In another embodiment, the processing unit 110 may analyze the geometry of the slow-moving trail by comparing the lateral distance traveled along the trail with the expected curvature of the slow-moving trail. The expected radius of curvature can be determined according to the following calculation: (δz2+δx2)/2/(δx), where δx represents the lateral distance traveled, and δz represents the longitudinal distance traveled. If the difference between the lateral distance traveled and the expected curvature exceeds a predetermined threshold (for example, 500 meters to 700 meters), the processing unit 110 may determine that the vehicle in front may change lanes. In another embodiment, the processing unit 110 may analyze the position of the preceding vehicle. If the position of the preceding vehicle conceals the road polynomial (for example, the preceding vehicle covers the top of the road polynomial), the processing unit 110 may determine that the preceding vehicle may change lanes. In the case where the position of the preceding vehicle is such that another vehicle is detected in front of the preceding vehicle and the slow-moving traces of the two vehicles are not parallel, the processing unit 110 may determine that the (closer) preceding vehicle may change lanes.

[0195]

In step 584, the processing unit 110 may determine whether the preceding vehicle 200 changes lanes based on the analysis performed in step 582.

For example, the processing unit 110 may make the determination based on the weighted average of the various analyses performed in step 582. Under such a scheme, for example, the decision that the preceding vehicle may change lanes made by the processing unit 110 based on a specific type of analysis may be assigned the value "1" (and "0" indicates that it is determined that the preceding vehicle is unlikely to change lanes). The different analyses performed in step 582 may be

IPR2025-00943
Tesla EX1002 Page 1203

assigned different weights, and the disclosed embodiments are not limited to any particular combination of analyses and weights.

[0196]

FIG. 6 is a flowchart illustrating an exemplary process 600 for generating one or more navigation responses based on stereoscopic image analysis consistent with the disclosed embodiments.

In step 610, the processing unit 110 may receive the first and second plurality of images via the data interface 128. For example, a camera included in the image acquisition unit 120 (such as image capture devices 122 and 124 with fields of view 202 and 204) may capture the first and second multiple images of the area in front of the vehicle 200 and be digitally connected (e.g., USB, wireless, Bluetooth, etc.) send the first and second plurality of images to the processing unit 110. In some embodiments, the processing unit 110 may receive the first and second plurality of images via two or more data interfaces. The disclosed embodiments are not limited to any specific data interface structure or protocol.

[0197]

In step 620, the processing unit 110 can execute the stereo image analysis module 404 to perform stereo image analysis of the first and second plurality of images to create a three-dimensional map of the road ahead of the vehicle and detect features in the image, such as lane markings, vehicles, Pedestrians, road signs, highway exit ramps, traffic lights, road hazards, etc.

The stereoscopic image analysis can be performed in a manner similar to the steps described above in conjunction with FIGS. 5A to 5D. For example, the processing unit 110 may execute the stereoscopic image analysis module 404 to detect candidate objects (for example, vehicles, pedestrians, road signs, traffic lights, road hazards, etc.) in the first and second plurality of images, based on various standards. Filter out a subset of candidates, perform multi-frame analysis, construct measurement results, and determine the confidence level of the remaining candidates. When performing the above steps, the processing unit 110 may consider information from both the first and second plurality of images, rather than information from a set of images alone. For example, the processing unit 110 may analyze differences in pixel level data (or other subsets of data from the two captured image streams) for candidate objects appearing in both the first and second plurality of images. As another example, the processing unit 110 may estimate the candidate object by observing that the object appears in one of the multiple images but does not appear in the other or other differences that may exist relative to the objects appearing in the two image streams. (E.g., relative to vehicle 200) position and/or speed. For example, the position, speed, and/or acceleration relative to the vehicle 200 may be determined based on the trajectory, position, movement characteristics, etc. of features associated with objects appearing in one or both of the image streams.

[0198]

In step 630, the processing unit 110 may execute the navigation response module 408 based on the analysis performed in step 620 and the technique described above in conjunction with FIG. 4 to generate one or more navigation responses in the vehicle 200.

For example, the navigation response may include turns, lane shifts, acceleration changes, speed changes, braking, etc. In some embodiments, the processing unit 110 may use data obtained from the execution speed and acceleration module 406 to generate one or more navigation responses. In addition, multiple navigation responses can occur simultaneously, sequentially, or in any combination thereof.

[0199]

FIG. 7 is a flowchart illustrating an exemplary process 700 for generating one or more navigation responses based on the analysis of three sets of images consistent with the disclosed embodiments.

In step 710, the processing unit 110 may receive the first, second, and third plurality of images via the data interface 128. For example, a camera included in the image capturing unit 120 (such as image capturing devices 122, 124, and 126 having fields of view 202, 204, and 206) can capture the first, second, and second regions of the front and/or side of the vehicle 200. Three more images and the first, second, and third plurality of images are sent to the processing unit 110 via a digital connection (for example, USB, wireless, Bluetooth, etc.). In some embodiments, the processing unit 110 may receive the first, second, and third multiple images via more than three data interfaces. For example, each of the image capturing devices 122, 124, 126 may have

an associated data interface for communicating data with the processing unit 110. The disclosed embodiments are not limited to any specific data interface structure or protocol.

[0200]

In step 720, the processing unit 110 may analyze the first, second, and third images to detect features in the images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, road hazards, etc. .

This analysis can be performed in a manner similar to the steps described above in conjunction with FIGS. 5A to 5D and FIG. 6. For example, the processing unit 110 may perform monocular image analysis on each of the first, second, and third plurality of images (for example, via executing the monocular image analysis module 402 and based on the steps described above in conjunction with FIGS. 5A to 5D). Alternatively, the processing unit 110 may perform stereoscopic image analysis on the first and second plurality of images, the second and third plurality of images, and/or the first and third plurality of images (for example, via the execution of the stereoscopic image analysis module 404 And based on the steps described above in conjunction with Figure 6). The analyzed processed information corresponding to the first, second, and/or third plurality of images can be combined. In some embodiments, the processing unit 110 may perform a combination of monocular image analysis and stereo image analysis. For example, the processing unit 110 may perform monocular image analysis on the first plurality of images (for example, via the execution of the monocular image analysis module 402) and perform stereoscopic image analysis on the second and third plurality of images (for example, via the execution of the stereoscopic image analysis module). 404). The configuration of the image capture devices 122, 124, and 126 (including their respective locations and fields of view 202, 204, and 206) can affect the type of analysis performed on the first, second, and third plurality of images. The disclosed embodiments are not limited to specifically configured image capture devices 122, 124, and 126, or the types of analysis performed on the first, second, and third plurality of images.

[0201]

In some embodiments, the processing unit 110 may perform a test of the system 100 based on the images acquired and analyzed in steps 710 and 720.

Such a test may provide an indicator of the overall performance of the system 100 to image capture devices 122, 124, and 126 of certain configurations. For example, the processing unit 110 may determine the ratio of "wrong hits" (for example, the system 100 incorrectly determines that there is a vehicle or pedestrian) and the ratio of "missing hits."

[0202]

In step 730, the processing unit 110 may generate one or more navigation responses in the vehicle 200 based on information obtained from both of the first, second, and third plurality of images.

The selection of both of the first, second, and third plurality of images may depend on various factors, for example, the number, type, and size of objects detected in each of the plurality of images. The processing unit 110 may also make selections based on the following: image quality and resolution, the effective field of view reflected in the image, the number of captured frames, the extent to which one or more objects of interest actually appear in the frame (for example, the object appears The percentage of the frame, the proportion of objects appearing in each such frame, etc.), etc.

[0203]

In some embodiments, the processing unit 110 may select the information obtained from two of the first, second, and third images by determining the degree of consistency between the information obtained from one image source and the information obtained from other image sources. information.

For example, the processing unit 110 may combine the processed information obtained from each of the image capturing devices 122, 124, and 126 (whether through monocular analysis, stereo analysis, or any combination of the two) and determine whether to span from the image capturing devices 122, 124 Each captured image in 126 presents consistent visual indicators (eg, lane markings, detected vehicles and their locations and/or paths, detected traffic lights, etc.). The processing unit 110 may also exclude information that is inconsistent across the captured image (for example, a vehicle that changes lanes, a lane model of a vehicle indicating that it is too close to the vehicle 200, etc.). Thus, the processing unit 110 can select information obtained from two of the first, second, and third plurality of images based on the determination of the coincidence and inconsistency information.

Tesla EX1002 Page 1205 IPR2025-00943

[0204]

For example, the navigation response may include turns, lane shifts, acceleration changes, and so on.

The processing unit 110 may generate one or more navigation responses based on the analysis performed in step 720 and the technique described above in conjunction with FIG. 4. The processing unit 110 may also use data obtained from the execution speed and acceleration module 406 to generate one or more navigation responses. In some embodiments, the processing unit 110 may generate one or more objects based on the relative position, relative velocity, and/or relative acceleration between the vehicle 200 and the detected object in any of the first, second, and third plurality of images. Navigation response. Multiple navigation responses can occur simultaneously, sequentially, or in any combination thereof.

[0205]

Smart lane offset

[0206]

The system 100 may provide driver assistance functionality to monitor and adjust the free space between the vehicle 200 and the travel boundary (eg, lane boundary) of the vehicle 200.

A lane may refer to a designated or expected path of travel of a vehicle and may have marked boundaries (e.g., lines on a road) or unmarked boundaries (e.g., road curbs, road barriers, guardrails, parked vehicles, etc.).

For example, by default, the system 100 may maximize or increase the free space between the vehicle 200 and the current lane boundary. Further, under certain conditions, the system 100 can cause the vehicle 200 to shift toward the side of the current travel lane without leaving the lane. For example, when the vehicle 200 is traveling around a curve, the system 100 can reduce the deviation toward the inside of the curve (for example, adjust the position of the vehicle 200 to be closer to the inside of the curve). As another example, in the case where an obstacle (such as a parked car, pedestrian, or cyclist) exists on one side of the lane, the system 100 may add an offset on the side of the vehicle 200 where the obstacle exists (for example, adjust the vehicle The position of 200 makes it farther away from obstacles).

[0207]

Figure 8A illustrates a vehicle 200 traveling on a road 800 in which the disclosed system and method for identifying lane constraints and operating the vehicle 200 within the lane constraints may be used.

As shown, the vehicle 200 may have a first vehicle side 802, and the first vehicle side 802 may be a first distance 805 from the first lane constraint 810. Similarly, the vehicle 200 may have a second vehicle side 812 opposite the first vehicle side 802, and the second vehicle side 812 may be a second distance 815 from the second lane constraint 820. In this way, the first lane constraint 810 and the second lane constraint 820 may define a lane in which the vehicle 200 can travel.

[0208]

The processing unit 110 may be configured to determine the first lane constraint 810 and the second lane constraint 820 based on a plurality of images acquired by the image capturing devices 122-126 that the processing unit 110 may receive via the data interface 128.

According to some embodiments, the first lane constraint 810 and/or the second lane constraint 820 may be identified by visible lane boundaries (such as dashed or solid lines marked on the road surface). Additionally or alternatively, the first lane constraint 810 and/or the second lane constraint 820 may include curbs or obstacles on the road surface. Additionally or alternatively, the first lane constraint 810 and/or the second lane constraint 820 may include markers (e.g., Bots point). According to some embodiments, the processing unit 110 may determine the first lane constraint 810 and/or the second lane constraint 820 by identifying the midpoint of the road surface width. According to some embodiments, if the processing unit 110 recognizes only the first lane constraint 810, the processing unit 110 may estimate the second lane constraint 820, such as based on the estimated lane width or road width. When, for example, the line specifying the lane of the road has not been painted or otherwise labeled, the processing unit 110 can recognize lane constraints in this manner.

[0209]

Tesla EX1002 Page 1206 IPR2025-00943

Detecting the first lane constraint 810 and/or the second lane constraint 820 may include: the processing unit 110 determines its three-dimensional model in the camera coordinate system.

For example, the three-dimensional model of the first lane constraint 810 and/or the second lane constraint 820 may be described by a third polynomial. In addition to three-dimensional modeling of the first lane constraint 810 and/or the second lane constraint 820, the processing unit 110 may perform multi-frame estimation of host motion parameters (such as the speed, yaw and pitch rate, and acceleration of the vehicle 200). Optionally, the processing unit 110 may detect static and moving vehicles and their positions, headings, speeds and accelerations (all relative to the vehicle 200). The processing unit 110 may further determine a road elevation model to convert all the information acquired from the multiple images into a three-dimensional space.

[0210]

Generally, as a default condition, the vehicle 200 may travel relatively centered within the first lane constraint and the second lane constraint 810 and/or 820.

However, in some situations, environmental factors can make it undesirable or unfeasible, such as when objects or structures exist on one side of the road. Thus, there may be situations where the vehicle 200 may desirably or advantageously be closer to one lane constraint or the other. The processing unit 110 may also be configured to determine whether such a situation referred to as a lane shift condition exists based on multiple images.

[0211]

In order to determine whether the lane deviation condition exists, the processing unit 110 may be configured to determine that one or more objects 825 exist in the vicinity of the vehicle 200.

For example, the object 825 may include another vehicle traveling on the road 800, such as a car, truck, or motorcycle. For each object 825, the processing unit 110 may determine an offset contour. Offset contour may include determining whether the object 825 will be located within a predetermined range of the vehicle 200 in its current predicted position. If the processing unit 110 determines that the object 825 is located or will be located within the predetermined range of the vehicle 200, the processing unit 110 may determine whether there is enough in the first lane constraint 810 and the second lane constraint 820 in order for the vehicle 200 to bypass the object 825. Space. If there is not enough space, the processing unit 110 may perform a lane change. If there is enough space, the processing unit 110 may determine whether there is enough time to bypass the object 825 before the distance between the vehicle 200 and the object 825 approaches. If there is enough time, the processing unit 110 may initiate or activate the offset maneuver to bypass the object 825. The offset maneuver may be performed by the processing unit 110 so that the movement of the vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 m/sec to 0.75 m/sec. The offset maneuver may be performed such that the maximum amplitude of the offset maneuver will occur when the distance between the vehicle 200 and the object 825 is shortened. If there is not enough time, the processing unit 110 may initiate braking (e.g., by sending an electronic signal to the braking system 230), and then initiate the excursion maneuver.

[0212]

After setting each offset contour for each object 825, the processing unit may combine the offset contours so that the maximum lateral distance is maintained with each object 825 bypassed by the vehicle 200.

That is, the offset maneuver may be modified and/or performed such that the lateral distance between the vehicle 200 and each object 825 bypassed by the vehicle 200 is maximized.

[0213]

Additionally or alternatively, the object 825 may be parked or stationary vehicles, walls, or people (such as pedestrians), or bicycle traffic.

The processing unit 110 may determine that a lane offset condition exists based on certain characteristics of the object 825. For example, the processing unit 110 may determine that if the height of the object 825 exceeds a predetermined threshold, such as

10 cm from the road surface, the object 825 constitutes a lane offset condition. In this way, the processing unit 110 may be configured to filter out objects for which lane deviation conditions will not occur, such as small pieces of debris.

[0214]

Additionally or alternatively, the processing unit 110 may be configured to determine whether the lane offset condition exists on the first vehicle side 802 based on the position of the target object 830 on the first side 802 of the vehicle.

The target object 830 may include one or more objects constituting the lane deviation condition. Additionally or alternatively, the target object 830 may not constitute a lane departure condition. For example, the target object 830 may be a street sign or other object that is far enough away from the lane of travel or that does not occur frequently enough (not constituting a lane departure condition). The processing unit 110 may be configured to conclude that the lane deviation condition does not exist if such a target object 830 is detected.

[0215]

The processing unit 110 may be configured to cause the vehicle 200 to travel such that the first distance 805 and the second distance 810 are substantially different to avoid or keep further away from the object 825.

For example, if the lane offset condition exists on the first vehicle side 802, the processing unit 110 may be configured to cause the vehicle 200 to travel within the first lane constraint 810 and the second lane constraint 820, such that the first distance 805 is greater than the second distance 815 . This is illustrated in Figure 8A, for example. Alternatively, if the lane offset condition exists on the second vehicle side 812, the processing unit 110 may create or execute such an offset profile configured to cause the vehicle 200 to constrain 810 and the second lane in the first lane. Travel within the lane constraint 820 such that the first distance 805 is smaller than the second distance 815. This is illustrated in Figure 8B, for example.

[0216]

As another example, a parked vehicle may be located on one side of the vehicle 200, and pedestrians may be located on the other side of the vehicle 200.

In such a situation, the system 100 may determine that it is desirable to reduce the distance between the vehicle 200 and the parked vehicle in order to provide more space between the vehicle 200 and the pedestrian. Therefore, the processing unit 110 may cause the vehicle 200 to travel so that the distance between parked vehicles is reduced.

[0217]

The recognition of the first vehicle side 802 and the second vehicle side 812 needs not to be static, so that at different points in time, the processing unit 110 can recognize the first vehicle side 802 as the driver side or the passenger side of the vehicle 200.

For example, depending on which side of the vehicle 200 one or more objects 825 are present, the recognition of the first vehicle side 802 may be changed so that if there is a lane offset condition based on the object 825, the object 825 is present on the first vehicle side 802 on. In such a situation, the processing unit 110 may create or execute an offset profile that is configured to cause the vehicle 200 to travel within the first lane constraint 810 and the second lane constraint 820 such that it is greater than the second distance 815 The first distance 805 of is provided towards the object 825.

[0218]

According to some embodiments, when a lane offset condition exists on the first vehicle side 802, the processing unit 110 may create an offset profile configured to cause the vehicle 200 to constrain 810 and the second lane in the first lane Travel within the constraint 820 such that the first distance 805 is at least 1.25 times greater than the second distance 815.

Additionally or alternatively, when a lane offset condition exists on the first vehicle side 802, the processing unit 110 may create an offset profile configured to cause the vehicle 200 to constrain 810 and the second lane in the first lane. Traveling within the two-lane constraint 820 makes the first distance 805 greater than the second distance 815 by at least 2 times.

[0219]

IPR2025-00943
Tesla EX1002 Page 1208

In some embodiments, as shown in FIG. 8C, the processing unit 110 may determine that the lane offset condition exists on both the first vehicle side 802 and the second vehicle side 812.

Then, the processing unit 110 may create such an offset profile for the offset condition that is configured to cause the vehicle 200 to travel within the first lane constraint 810 and the second lane constraint 820 such that the first distance 805 is substantially It is equal to the second distance 815.

[0220]

The processing unit 110 may be configured to prioritize the offset contours for the moving object 825 and the stationary object 825.

For example, the processing unit 110 may preferentially offset the contour to maintain a predetermined distance between the vehicle 200 and the moving object 825 and/or the static object 825. The processing unit 200 may prioritize these offset conditions (better than these offset conditions) based on certain lane characteristics (such as a curve). The last priority lane departure condition may include imitating the preceding vehicle.

[0221]

The processing unit 110 may condition the offset profile or offset into a desired look-ahead point, which may be located in the middle of the lane defined by the first lane constraint 810 and the second lane constraint 820.

The processing unit 110 may calculate and implement the yaw rate command based on the look-ahead point and the offset profile. The yaw rate command may be used to calculate the torque command based on the yaw rate command and the yaw rate of the vehicle 200 measured by using a gyroscope. The processing unit 110 may send a torque command to the steering system of the vehicle 200 (for example, the steering system 240).

[0222]

According to some embodiments, such as via the user interface 170, an input from the driver of the vehicle 200 may cancel the processing unit 110.

For example, a user input can cancel the control of the processing unit 110, causing the vehicle 200 to travel within the first lane constraint 810 and the second lane constraint 820, so that the first distance 805 is greater than the second distance 815.

[0223]

Additionally or alternatively, the processing unit 110 may be configured to determine whether the lane deviation condition is substantially non-existent based on a plurality of images captured by the image capturing device 122.

This may include situations in which one or more objects 825 below a certain size threshold are detected, and/or when no object 825 is detected, the detected one or more objects 825 are located at a predetermined distance from the lane constraint Outside. In such a situation, the processing unit 110 may be configured to cause the vehicle 200 to travel within the first lane constraint 810 and the second lane constraint 820 such that the first distance 805 is equal to or substantially equal to the second distance 815.

[0224]

As illustrated in FIG. 8D, the lane deviation condition may include a curved road.

The processing unit 110 may be configured to cause the vehicle 200 to travel within the first lane constraint 810 and the second lane constraint 820, such that the first distance 805 is toward the outside of the curve relative to the vehicle 200, and the second distance 815 is toward the curve relative to the vehicle 200. Inside. For example, the first distance 805 may be smaller than the second distance 815.

[0225]

FIG. 9 is an exemplary block diagram of the memory 140 and/or 150 consistent with the disclosed embodiments. The memory 140 and/or 150 may store instructions for performing one or more operations.

As illustrated in FIG. 9, the memory 140 may store one or more modules for performing the lane deviation condition detection and response described herein. For example, the memory 140 may store the lane restriction module 910 and the lane offset module 920. The application processor 180 and/or the image processor 190 can execute instructions stored in any of the modules 910 and 920 (included in the memory 140). Those skilled in the art will understand that the references to the processing unit 110 in the following discussion may refer to the application processor 180 and the image processor 190 individually or collectively. Therefore, the steps of any of the following processes may be executed by one or more processing devices.

[0226]

The lane restriction module 910 may store instructions, and when executed by the processing unit 110, the instructions may detect and define the first lane restriction 810 and the second lane restriction 820.

For example, the lane offset module 910 may process multiple images received from at least one image capture device 122-124 to detect the first lane constraint 810 and the second lane constraint 820. As discussed above, this may include identifying painted lane lines and/or measuring the midpoint of the road surface.

[0227]

The lane offset module 920 may store instructions. When executed by the processing unit 110, the instructions may detect the existence of lane offset conditions and/or identify whether the lane offset conditions exist on the first vehicle side 802 and/or the second vehicle side. 812 on.

For example, the lane offset module 920 may process multiple images to detect the presence of the object 825 on the first vehicle side 802 and/or the second vehicle side 812. The lane offset module 920 may process multiple images to detect a curve in the road on which the vehicle 200 is traveling. Additionally or alternatively, the lane offset module 920 may receive information from another module or other system that indicates the presence of an object 825 and/or a curved lane. The lane offset module 920 may perform control to change the position of the vehicle 200 relative to the first lane constraint 810 and/or the second lane constraint 820. For example, if the lane offset module 920 determines that the lane offset condition is only on the first vehicle side 802, the lane offset module 920 may cause the vehicle 200 to move closer to the second lane constraint 820. On the other hand, if the lane offset module 920 determines that the lane offset condition is only on the first vehicle side 812, the lane offset module 920 may cause the vehicle 200 to move closer to the first lane constraint 810.

[0228]

FIG. 10 illustrates a process 1000 for navigating the vehicle 200 consistent with the disclosed embodiment.

The process 1000 may identify lane constraints that define the lane of travel in which the vehicle travels. The process 1000 includes causing the vehicle 200 to travel closer to one or another lane constraint, such as the first lane constraint 810 or the second lane constraint 820 in response to detecting a lane offset condition. In this way, the process 1000 can be used to drive the vehicle 200 within lane constraints and to consider environmental factors that may be more desirable to displace the vehicle 200 within the first lane constraint 810 and the second lane constraint 820.

[0229]

At step 1010, the process 1000 may include using at least one image capture device 122, 124, and/or 126 to acquire multiple images of the area near the vehicle 200.

For example, the processing unit 110 may receive multiple images via the data interface 128. For example, the processing unit 110 may be configured to determine that one or more objects 825 exist on one side of the road. For example, the object 825 may be a parked or stationary vehicle, wall or person (such as pedestrian) or bicycle traffic. The processing unit 110 may determine whether the lane deviation condition exists based on certain characteristics of the object 825. For example, the processing unit 110 may determine that if the height of the object 825 exceeds a predetermined threshold, such as 10 cm from the road surface, the object 825 constitutes a lane offset condition. In this way, the processing unit 110 may be configured to filter out objects for which lane deviation conditions will not occur, such as small pieces of debris.

[0230]

Additionally or alternatively, the processing unit 110 may be configured to determine whether the lane offset condition exists on the first vehicle side 802 based on the position of the target object 830 on the first side 802 of the vehicle.

The target object 830 may be one or more objects constituting the lane deviation condition. Additionally or alternatively, the target object 830 may not constitute a lane departure condition. For example, the target object 830 may be a road obstacle or a street sign. The processing unit 110 may be configured to conclude that the lane offset condition does not exist if such a target object 830 is detected.

[0231]

In step 1020, the process 1000 may determine the first lane constraint on the first vehicle side 802 and the second lane constraint on the second vehicle side 812 based on the multiple images.

For example, the processing unit 110 may be configured to determine the first lane constraint 810 and the second lane constraint 820 based on a plurality of images received via the data interface 128.

[0232]

At step 1030, the process 1000 may determine whether a lane offset condition exists on the first vehicle side 802.

If the lane offset condition does exist on the first vehicle side 802, then at step 1040, the process 1000 may include: causing the vehicle 200 to travel within the first lane constraint 810 and the second lane constraint 820, so that the second distance 815 is less than the first A distance of 805. The processing unit 110 may be configured to cause the vehicle 200 to travel such that the first distance 805 and the second distance 810 are substantially different to avoid or keep further away from the object 825. For example, if the lane offset condition exists on the first vehicle side 802, the processing unit 110 may be configured to cause the vehicle 200 to travel within the first lane constraint 810 and the second lane constraint 820, such that the first distance 805 is greater than the second distance 815. This is illustrated in Figure 8A, for example. Alternatively, if the lane offset condition exists on the second vehicle side 812, the processing unit 110 may be configured to cause the vehicle 200 to travel within the first lane constraint 810 and the second lane constraint 820, so that the first distance 805 is less than the second The distance is 815. This is illustrated in Figure 8B, for example. For example, in order to navigate the vehicle 200 according to the process 1000, the processing unit 110 may send electronic signals to one or more of the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200.

[0233]

If the lane offset condition does not exist on the first vehicle side 802, then in step 1050, the method may include: determining whether the lane offset condition exists on the second vehicle side 812.

If the lane offset condition exists on the second vehicle side 812, the method may include causing the vehicle 200 to travel between the first lane constraint 810 and the second lane constraint 820 such that the first distance 805 is less than the second distance 815.

[0234]

Navigate the car to the default lane

[0235]

The system 100 may provide driver assistance functionality for monitoring the location of the vehicle 200 in the current travel lane and if the current travel lane is not the predetermined default travel lane, the vehicle 200 enters the predetermined default travel lane.

The predetermined default lane of travel may be defined by user input and/or as a result of processing images received from at least one image capture device 122-126.

[0236]

The system 100 may also provide the driver and passengers of the vehicle 200 and/or other drivers with a notification that the vehicle 200 is about to change lanes to a predetermined default lane of travel.

For example, the system 100 may activate a turn signal located at or near one of the bumper areas 210.

Additionally or alternatively, the system 100 may issue a voice notification, such as through a speaker 360. When the vehicle is traveling, the vehicle 200 may use the system 100 to select a specific lane as a predetermined default traveling lane. When the current lane in which the vehicle 200 travels (for example, the current travel lane) is different from the predetermined default travel lane, the system 100 may cause the vehicle 200 to make a lane change, for example, enter the predetermined default travel lane.

[0237]

FIG. 11 illustrates a vehicle 200 traveling in the current traveling lane 1110 of the road 800.

According to some embodiments, the vehicle 200 may have a predetermined default lane 1120 of travel. For example, the predetermined default travel lane 1120 may be the rightmost lane of the road 800. The navigation system 100 may be configured to determine whether the vehicle 200 is traveling in a predetermined default travel lane 1120, and if not, cause the vehicle 200 to change to the predetermined default travel lane 1120. For example, the processing unit 110 may be configured to compare the current travel lane 1110 with a predetermined default travel lane 1120 and cause the vehicle 200 to return to the predetermined default travel lane 1120.

[0238]

According to some embodiments, the processing unit 110 may detect the object 1130 in the predetermined default travel lane 1120 based on a plurality of images acquired by one or more of the image capturing devices 122-126.

The processing unit 110 may detect the position and speed of the vehicle 200 relative to the object 1130 by analyzing the image using the technique described above in conjunction with FIGS. 5A to 5D. According to some embodiments, the processing unit 110 may determine whether the vehicle 200 should overtake the object 1130. It may be desirable, for example, that the speed and/or acceleration of the vehicle 200 relative to the object 1130 exceed a certain threshold.

[0239]

FIG. 12A is an exemplary block diagram of the memory 140 consistent with the disclosed embodiments. The memory 140 may store instructions for performing one or more operations.

As illustrated in FIG. 12A, the memory 140 may store one or more modules for performing the detection and response of the default travel lane and the current travel lane described herein. For example, the memory 140 may store the current lane detection module 1210, the default lane detection module 1220, and the action module 1230.

[0240]

The current lane detection module 1210 can store instructions, and when executed by the processing unit 110, the instructions can detect and limit the current lane 1110 of travel.

For example, the processing unit 110 may execute the current lane detection module 1212 to process a plurality of images received from at least one image capturing device 122-124 and detect the current traveling lane 1110. This may include detecting the relative position of the vehicle 200 relative to the curb of the road 800, the midpoint of the road 800, and/or lane markers (such as painted lane lines). The processing unit 110 may detect the position of the vehicle 200 relative to the road 800 by analyzing the image using the technique described above in conjunction with FIGS. 5A to 5D.

[0241]

The default lane detection module 1220 may store instructions, and when executed by the processing unit 110, the instructions may identify a predetermined default lane 1120 and compare the default lane detection module 1220 with the current lane 1110 identified by the current lane detection module 1210.

The default lane detection module 1220 may identify the predetermined default lane of travel 1120 based on user input and/or multiple images received from at least one image capture device 122-124. For example, the user may select the default lane through a voice command received by the speaker 360 or by selecting a menu displayed on the touch screen 320.

[0242]

In some embodiments, the predetermined default travel lane 1120 may be the rightmost travel lane, or may be a lane with more desirable characteristics.

For example, the default lane detection module 1220 may include instructions for selecting a lane with less traffic, the lane of the vehicle closest to the speed of the vehicle 200, and/or the lane closest to an exit or the street where the vehicle 200 will turn. . In some embodiments, the predetermined default lane of travel 1120 may be determined based on data received from the position sensor 130 (such as a global positioning system).

[0243]

The action module 1230 may store an instruction, and when executed by the processing unit 110, the instruction may generate a response identifying the current travel lane 1110 and a predetermined default travel lane 1120.

For example, the processing unit 110 may execute the action module 1230 to cause the vehicle 200 to change lanes when the current travel lane 1110 is different from the predetermined default travel lane 1120. Additionally or alternatively, the action module 1230 may initiate a lane change notification, such as activating the turn signal 210 or generating a voice notification.

[0244]

The action module 1230 may include instructions that cause the processing unit 110 to determine whether the vehicle 200 should move out of the predetermined default lane 1120 of travel.

For example, the action module 1230 may determine that the object 1130 in the predetermined default travel lane 1120 should be bypassed by the vehicle 200.

[0245]

The action module 1230 may further include instructions that cause the processing unit 110 to determine whether it is safe to perform a lane change.

For example, before performing a lane change, the processing unit 110 may evaluate and determine one or more characteristics of the road 800 based on, for example, analysis of multiple images performed using the technique described above in conjunction with FIGS. 5A to 5D. If the analysis indicates that the lane change condition is unsafe, the processing unit 110 may determine that the vehicle 200 should remain in the current lane.

[0246]

The action module 1230 may include additional instructions for manipulating the vehicle 200 to maneuver out and return to the predetermined default travel lane 1120 (such as bypassing the object 1130).

When executing these instructions, the processing unit 110 may use the estimated motion of the vehicle 200 to estimate the location of the predetermined default lane 1120 of travel. This may be advantageous in situations where the predetermined default lane of travel 1120 cannot be kept as detected by the image capture devices 122-124.

[0247]

The processing unit 110 may be configured to execute an offset profile that offsets the vehicle 200 from a predetermined default travel lane 1120 by a full lane width, which may be estimated based on the distance between detected lane markers.

The offset maneuver defined by the offset profile may be performed by the processing unit 110 so that the movement of the vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 m/sec to 0.75 m/sec. During the offset maneuver, the processing unit 110 can execute an instruction from the action module 1230, and the instruction switches to follow the center of the current traveling lane 1110. These instructions may include a gradual transition

from following the predicted lane of the belt offset to following the current lane 1110 of travel. The gradual transition can be done by defining the desired look-ahead point so that no more changes in time than the recommended slew rate (which can be about 0.25 m/s) are allowed.

[0248]

The instructions for making the vehicle 200 transition from the predetermined default lane of travel 1120 may be used to make the vehicle 200 transition back into the predetermined default lane of travel.

For example, the processing unit 110 may be configured to execute an offset profile that offsets the vehicle 200 from the current travel lane 1110 by a full lane width, which may be estimated based on the distance between detected lane markers. The offset maneuver defined by the offset profile may be performed by the processing unit 110 so that the movement of the vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 m/sec to 0.75 m/sec. During the shifting maneuver, the processing unit 110 can execute an instruction from the action module 1230, and the instruction switches to follow the center of the predetermined default travel lane 1120. These instructions may include a gradual transition from following the predicted lane of the belt offset to following the predetermined default lane of travel 1120. The gradual transition can be done by defining the desired look-ahead point so that no more changes in time than the recommended slew rate (which can be about 0.25 m/s) are allowed.

[0249]

In one or more steps of making the vehicle 200 transition between two lanes, the processing unit 110 may calculate and implement a yaw rate command based on the look-ahead point and the offset profile.

The yaw rate command may be used to calculate the torque command based on the yaw rate command and the yaw rate of the vehicle 200 measured using a gyroscope. The processing unit 110 may send a torque command to the steering system of the vehicle 200 (for example, the steering system 240).

[0250]

FIG. 12B illustrates a process 1250 for navigating the vehicle 200 to the default lane consistent with the disclosed embodiment.

According to some embodiments, the process 1250 may be implemented by one or more components of the navigation system 100 (such as at least one processing unit 110). The process 1250 may identify the current lane 1110 of the vehicle 200 and compare it with the predetermined default lane 1120 to determine whether the vehicle 200 is traveling within the predetermined default lane 1120. If the vehicle 200 does not travel in the predetermined default travel lane 1120, the processing unit 110 may cause the vehicle 200 to change lanes so as to navigate to the default travel lane 1120.

[0251]

In step 1260, one or more of the image capturing devices 122-124 may acquire multiple images of the area near the vehicle 200.

The processing unit 110 may receive multiple images via the data interface 128.

[0252]

In step 1270, the processing unit 110 may execute the current lane detection module 1210 to detect and define the current lane 1110 of travel.

For example, the processing unit 110 may process multiple images, for example, to detect the relative position of the vehicle 200 with respect to the curb of the road 100, the midpoint of the road 100, and/or the painted lane markers. The processing unit 110 may detect the position of the vehicle 200 relative to the road 800 by analyzing the image using the technique described above in conjunction with FIGS. 5A to 5D.

[0253]

In step 1280, the processing unit 110 may execute the default lane detection module 1220 to identify a predetermined default lane 1120 and compare the default lane detection module 1220 with the current lane 1110 identified by the current lane detection module 1210.

The default lane detection module 1220 may identify the predetermined default lane of travel 1120 based on user input and/or multiple images received from at least one image capture device 122-124.

[0254]

For example, as discussed above, the processing unit 110 may determine the predetermined default lane of travel 1120 based on any number of factors.

The predetermined default travel lane 1120 may be the rightmost lane among the plurality of travel lanes 1110 and 1120. Additionally or alternatively, the determination of the predetermined default lane of travel 1120 may be based on input received from the user via the user interface 170 by processing a plurality of images, and/or based on data received from the position sensor 130 (such as a global positioning system). The predetermined default travel lane 1120 may be determined once for each vehicle trip, or may be determined periodically. For example, the predetermined default lane of travel 1120 may be dynamically determined in response to the conditions of a given location.

[0255]

If the current lane 1110 is not the same as the predetermined default travel lane 1120, the processing unit 110 may execute the action module 1230 to cause the vehicle 100 to navigate to the default travel lane 1120.

For example, the processing unit 110 may send an electronic signal to one or more of the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200 to trigger a desired response. For example, the processing unit 110 may cause the steering system 240 to steer the steering wheel of the vehicle 200 to achieve a predetermined angle of rotation.

[0256]

Further, before performing the lane change, the processing unit 110 may determine whether the lane change is safe (for example, there are no vehicles or objects on the road).

As discussed above, the processing unit 110 may evaluate one or more characteristics of the road 800 based on, for example, analysis of multiple images performed using the techniques described above in conjunction with FIGS. 5A to 5D. If the lane change condition is unsafe, the processing unit 110 may determine that the vehicle 200 should remain in the current lane.

[0257]

In some implementations, the execution action module 1230 can initiate a lane change notification.

For example, the processing unit 110 may generate a signal to activate a turn signal and/or generate a voice notification via the speaker 360. Therefore, such notification may occur before the vehicle 200 changes lanes. In other embodiments, the execution action module 1230 may initiate a notification to the driver of the vehicle 200 that the vehicle 200 is not traveling in the default lane without causing the lane change to occur. Therefore, in such an embodiment, the driver of the vehicle 200 may determine whether to manually take an action and steer the vehicle 200 to change lanes.

[0258]

In some embodiments, the system 100 may notify the user of the vehicle 200 and/or other drivers in the vicinity of the vehicle 200 that if the vehicle 200 will move to a predetermined default lane of travel 1120, the vehicle 200 will change lanes.

For example, the processing unit 110 may activate the turn signal of the vehicle 200 before causing the vehicle 200 to change lanes. Additionally or alternatively, a voice notification notifying other drivers in the vicinity may be generated by the processing unit 110. The voice notification may be sent to a receiver included in other nearby vehicles via, for example, Bluetooth. Additionally or alternatively, the at least one processing unit 110 may cause a voice indicator to notify the driver of the vehicle 200 before causing the vehicle 200 to switch lanes. In this way, the driver of the vehicle 200 can expect that the control system 100 will be applied to the vehicle 200.

IPR2025-00943
Tesla EX1002 Page 1215

[0259]

Control the speed of turning vehicles

[0260]

The system 100 may provide driver assistance functionality for controlling the speed (ie, speed and/or direction) of the vehicle 200 in different scenarios, such as when making a turn (eg, passing a curve).

For example, the system 100 may be configured to use a combination of map, location, and/or visual data to control the speed of the vehicle 200 when turning.

In particular, the system 100 can control the speed of the vehicle 200 when approaching a curve (while navigating the curve) and/or when leaving the curve. The system 100 may consider different factors to control the speed, depending on the position of the vehicle 200 relative to the curve at a given time. For example, the system 100 may initially cause the vehicle 200 to decelerate in response to learning that the curve is a certain distance in front of the vehicle 200. However, while the vehicle 200 is turning through a curve, the system 100 may cause the vehicle 200 to accelerate and/or change direction based on the analysis of the characteristics of the curve.

[0261]

13A and 13B are schematic representations of bringing a vehicle (for example, the vehicle 200) close to and navigating a curve having one or more characteristics of the road 1300 consistent with the disclosed embodiment.

As illustrated in FIG. 13A, a visual indicator (such as a road sign 1302 located near the road 1300) may inform the vehicle 200 that the curve is ahead. In response to the warning provided by the road sign 1302, the driver of the vehicle 200 and/or the system 100 that provides driver assistance functionality for the vehicle 200 may adjust the speed of the vehicle 200 (for example, by decelerating and/or steering) to safely Navigate corners.

[0262]

When the vehicle 200 is approaching a curve, the system 100 may detect the lane width 1306 and the curve radius 1308 by using the image capturing device 122 and/or 124 (which may be a camera) of the vehicle 200.

Based on the curve radius 1308, the lane width 1306, and other characteristics associated with the curve detected using the image capture device 122 and/or 124, the system 100 may adjust the speed of the vehicle 200. Specifically, the system 100 may adjust the speed of the vehicle 200 when the vehicle 200 approaches a curve and/or the vehicle 200 navigates the curve. The system 100 may also adjust the speed of the vehicle 200 as it leaves the curve in response to, for example, a traffic light 1304 that appears after the curve. The following describes further details on adjusting the speed of the vehicle 200 and detecting the lane width 1306, the radius of curvature or the radius of a curve 1308, and the characteristics associated with the curve in conjunction with FIGS. 14 and 15.

[0263]

In some embodiments, the system 100 may identify the curve to be navigated based on map data and/or vehicle location information (eg, GPS data).

The system 100 may determine the initial target speed of the vehicle based on one or more characteristics of the curve reflected in the map data. The system 100 may adjust the speed of the vehicle to the initial target speed, and determine one or more observed characteristics of the curve based on one or more images acquired by one or more of the image capture devices 122-126. The system 100 may determine the updated target speed based on the observed one or more characteristics of the curve, and adjust the speed of the vehicle to the updated target speed.

[0264]

FIG. 14 is an exemplary block diagram of a memory consistent with the disclosed embodiments, the memory being configured to store instructions for performing one or more operations.

As indicated in FIG. 14, the memory 140 may store a curve recognition module 1402, a characteristic observation module 1404, and a speed module 1406. The disclosed embodiments are not limited to the memory 140 of any particular configuration. Further, the processing unit 110 can execute instructions stored in any of the modules 1402-1406 (included in the memory 140).

[0265]

In one embodiment, the curve recognition module 1402 may store software instructions (such as computer vision software), and when executed by the processing unit 110, the software instructions can analyze one or more sets of images and/or use maps and/or Or location data (for example, data stored in the map database 160) to detect a curve.

The image group may be acquired by one or more of the image capture devices 122, 124, and 126. Based on this analysis, the system 100 (eg, via the processing unit 110) may cause the speed of the vehicle 200 to change. For example, if the processing unit 110 detects that there is a curve within the minimum threshold distance in front of the vehicle 200, the processing unit 110 may cause the speed of the vehicle 200 to decrease by a predetermined amount and/or adjust the direction of the vehicle 200 by a specific angle.

[0266]

In one embodiment, the characteristic observation module 1404 may store software instructions (such as computer vision software). When executed by the processing unit 110, the software instructions observe the curve detected by the system 100 (for example, by executing the curve recognition module 1402). Dao-related characteristics.

The processing unit 110 may analyze one or more sets of images acquired by one or more of the image capture devices 122, 124, and 126 and/or use map and/or location data (for example, stored in the map database 160). Data) to observe the characteristics associated with the detected curve. The analysis can generate information associated with the curve, such as the radius of the curve 1308, the width of the lane 1306, the degree of curvature, the rate of change of curvature, the inclination of the curve, the length or arc length of the curve, and so on. These characteristics may include estimates (for example, curve characteristics calculated using pre-existing data and/or mathematical models based on such data) and actual values (for example, curve characteristics calculated by analyzing captured images of the curve) . Based on this analysis, the processing unit 110 may cause the speed of the vehicle 200 to change. For example, in view of the higher rate of curvature change associated with a curve 25 meters in front of the vehicle 200, the processing unit 110 may cause the speed of the vehicle 200 to decrease and/or adjust the direction of the vehicle 200 by a specific angle.

[0267]

In one embodiment, the speed module 1406 may store software instructions configured to analyze data from one or more computerized electrical devices that are configured to determine the target speed and cause the speed of the vehicle 200 to change is the target speed.

For example, the processing unit 110 may execute the speed module 1406 to calculate the target speed of the vehicle 200 based on the data obtained from executing the curve recognition module 1402 and the characteristic observation module 1404. For example, such data may include initial target position and initial speed, updated target position and update speed, final target position and final speed, and so on. In addition, the processing unit 110 may calculate the speed of the vehicle 200 based on input from other systems of the vehicle 200 (such as the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200). Based on the calculated target speed, the processing unit 110 may send electronic signals (for example, via a controller area network bus ("CAN bus")) to the throttle system 220, braking system 230, and/or steering system 240 of the vehicle 200 to The change in the speed and/or direction of the vehicle 200 is triggered by, for example, physically depressing the brake, releasing the accelerator of the vehicle 200, or turning the vehicle 200 in a specific direction (for example, at a specific angle). The increase in speed may be associated with a corresponding acceleration (for example, between 0.2 m/s2 and 3.0 m/s2). Conversely, the speed reduction can be associated with the corresponding deceleration. Additionally, acceleration or deceleration may be based on road type (eg, highway, city street, country road, etc.), the presence of any nearby speed constraints (eg, sharp curves, traffic lights), and road width or lane width 1306. For the purposes of this disclosure, deceleration may refer to acceleration having a negative value.

[0268]

Tesla EX1002 Page 1217 IPR2025-00943

15 is a flowchart showing an exemplary process 1500 for controlling the vehicle speed based on the detected curve and the observed curve characteristics and/or map data about the curve, consistent with the disclosed embodiment.

In step 1510, the processing unit 110 may receive a plurality of images via the data interface 128 between the processing unit 110 and the image acquisition unit 120. For example, a camera (such as the image capturing device 122) included in the image acquisition unit 120 may capture a plurality of images and send the images to the processing unit 110 via a digital connection (eg, USB, wireless, Bluetooth, etc.). In some embodiments, the processing unit 110 may receive more than one multiple images via multiple data interfaces. For example, the processing unit 110 may receive multiple images from each of the image capture devices 122, 124, 126, and each of the image capture devices 122, 124, 126 may have an associated data interface for communicating with the processing unit 110 data. The disclosed embodiments are not limited to any specific data interface structure or protocol.

[0269]

In step 1520, the processing unit 110 may execute the curve recognition module 1502 to detect a curve located in front of the vehicle 200 based on, for example, map and/or location data (for example, data stored in the map database 160).

Such data can indicate the location of a curve that exists on a given road. The data may indicate the absolute position of a particular curve (for example, via GPS coordinates) or the relative position of the curve (for example, by describing the curve relative to the vehicle 200, road signs 1302, traffic lights 1304, and/or other descriptive landmarks) . In some embodiments, the processing unit 110 may execute the curve recognition module 1502 to detect a curve by analyzing a plurality of images. This analysis can be performed using the technique described above in conjunction with FIGS. 5A to 5D and FIG. 6.

[0270]

In step 1530, the processing unit 110 may execute the speed module 1406 to determine the initial target speed of the vehicle 200 based on, for example, data obtained from executing the curve recognition module 1402.

The initial target speed may reflect the speed and direction of the vehicle 200 to safely enter the upcoming curve, and may be determined using the technique described above in conjunction with FIG. 14. For example, the speed module 1406 may determine the initial target speed based on known or estimated characteristics associated with the upcoming curve, such as the curve radius 1308 and the inclination of the curve. The speed module 1406 may calculate the initial target speed according to mathematical models and/or formulas and based on various constraints. In some embodiments, for example, the speed module 1406 may be based on the presence of a traffic light (such as a traffic light 1304) before, during, or after a curve, weather conditions (e.g., light rain, heavy rain, strong wind, cloudy, etc.), road conditions (For example, asphalt roads, dirt roads, potholes, etc.), speed limit, nearby vehicles or pedestrians, etc. to determine the initial target speed. The initial target speed may also be based on the lateral acceleration limit associated with the vehicle 200. In some embodiments, the lateral acceleration limit may be in the range of 0.2 m/sec 2 to 3.0 m/sec 2 and may be adjusted depending on the curve characteristics, the width of the road, or the width of the lane 1306. Additionally, the initial target speed may be based on the focal range associated with one or more cameras (such as image capture devices 122 and 124).

[0271]

For example, the speed module 1406 may determine the initial target speed based on the radius of curvature of the road in front of the vehicle 200.

The speed module 1406 may determine the radius of curvature at various distances from the vehicle 200 and store the information in a vector R(1:n). In one embodiment, the speed module 1406 may determine the radius of curvature of the following locations: (1) the current position of the vehicle 200; (2) the distance from the vehicle 200 D end point, where D end point represents the recognition of the road characteristics before the vehicle 200 The distance that allows the system 100 to control the navigation response of the vehicle 200 so that the driver can feel comfortable; and (3) the incremental distance from the vehicle 200, such as every 1 meter, and so on. The value of the distance D end point may depend on the current speed Vego of the vehicle 200. For example, the greater the value of Vego (e.g., the faster the vehicle 200 travels), the greater the value of D end point (e.g., the longer the distance can be considered to provide enough time for the system 100 to provide a navigation response, such as based on approaching a sharp turn Or curve to reduce the speed of the vehicle 200 comfortably). In one embodiment, the D endpoint can be calculated according to the following formula: D endpoint=(Vego2/2)/(0.7/40*Vego+0.5). Thus, the vector D(1:n)=[1:1:n], where n=round down (D end point).

[0272]

In some embodiments, the speed module 1406 may determine the initial target speed based on the lateral acceleration constraint amax (eg, the maximum allowable lateral acceleration).

a can fall within the range of 1.5 m/s2 to 3.5 m/s2 at most. Based on a maximum and the radius of curvature vector R (discussed above), the speed module 1406 may limit the speed of the vehicle 200 (for example, at different points or distances from the current position of the vehicle 200) to the speed v limit, calculated as: v Limit (1:n)=(amax*R(1:n))1/2. In addition, the deceleration constraint derived from the determination of the v limit can be calculated as: dec limit=minimum{((vlimit(1:n) 2-Vego2)/2), maximum{dmin, D(1:n) -2*Vego}}, where d is at least equal to 40 meters. The smallest value of d may indicate a point at which the driver of the vehicle 200 feels comfortable to decelerate.

[0273]

In some embodiments, the speed module 1406 may determine the initial target speed based on the curvature slew rate constraint slew maximum (which may be based on the rate of curvature change).

The maximum slew can fall within the range of 0.5 m/s3 to 1.5 m/s3. The speed module 1406 can determine the v limit according to the following formula: v limit (1:n)=(slew max/abs{{1/R(1:n)}})1/3. Further, the speed module 1406 can determine the dec limit according to the following formula: dec limit=minimum{((vlimit(1:n)2-Vego2)/2), maximum{dminimum, D(1:n)-2* Vego}}, where d is at least equal to 40 meters.

[0274]

In some embodiments, the speed module 1406 may determine comfort constraints such that the driver of the vehicle 200 is comfortable while the vehicle 200 experiences various navigation responses, such as turning or navigation curves.

The speed module 1406 may determine the v limit according to the following formula: v limit (1:n)=4/500*R(1:n)+22.5. Further, the speed module 1406 can determine the dec limit according to the following formula: dec limit=minimum{((vlimit(1:n) 2-Vego2)/2), maximum{dminimum, D(1:n)-2* Vego}}, where d is at least equal to 40 meters.

[0275]

The speed module 1406 may merge the three constraints determined above, namely the lateral acceleration constraint, the curvature slew rate constraint, and the comfort constraint by using the minimum value of the dec limit calculated in association with each constraint.

After merging the constraints, the speed module 1406 can reach the initial target speed by determining the amount to adjust the speed of the vehicle 200 according to the following formula, denoted as v comfort: v comfort = v comfort (previous) + dec limit * Δt, where v comfort (Previous) represents the previous speed of the vehicle 200 (eg, the previously determined speed, the speed of the vehicle 200 before any adjustments are made by the speed module 1406, etc.), and Δt represents the time between the determination of v comfort and v comfort (previous) .

[0276]

In step 1540, the processing unit 110 may execute the speed module 1406 to adjust the speed of the vehicle 200 to the initial target speed determined in step 1530.

Based on the current speed of the vehicle 200 and the initial target speed determined in step 1530, the processing unit 110 may send electronic signals to the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200 to change the speed of the vehicle 200 To the initial target speed. Further, one or more actuators may control the throttle system 220, the braking system 230, and/or the steering system 240. For example, the processing unit 110 may send an electronic signal to cause the system 100 to physically depress the brake for a predetermined amount or partially release the accelerator of the vehicle 200. Further, the processing unit 110 may send an electronic signal to cause the system 100 to steer the vehicle 200 in a specific direction. In some embodiments, the magnitude of the acceleration (or deceleration) may be based on the difference between the current speed of the vehicle 200 and the initial target speed. When the difference is large (for example, the current speed is 10km/hr greater than the initial target speed), for example, the system 100 can maximize the deceleration and cause the vehicle 200 to achieve the initial target as safely as possible in the shortest amount of time. Speed mode brake to

change the current speed. Alternatively, the system 100 may apply the brakes in a manner that minimizes the deceleration and causes the vehicle 200 to gradually achieve the initial target speed. More generally, the system 100 can change the current speed to the initial target speed according to any particular braking profile (eg, a braking profile that requires a high braking level for the first two seconds and a low braking level for the following three seconds). The disclosed embodiments are not limited to any specific braking profile or braking method.

[0277]

In step 1550, the processing unit 110 may execute the characteristic observation module 1404 to determine one or more characteristics of the curve located in front of the vehicle 200 based on, for example, analyzing multiple images performed using the technique described above in conjunction with FIGS. 5A to 5D.

For example, the characteristics may include estimated and/or actual values associated with the radius of the curve 1308, the width of the road or lane 1306, the degree of curvature, the rate of change of curvature, the inclination of the turn, and the like. For example, the characteristics may also include information conveyed by one or more road signs, such as road signs 1302. In addition, the characteristic may include the length or arc length of the detected curve (which may be a function of the speed of the vehicle 200 when approaching the curve). The processing unit 110 may execute the characteristic observation module 1404 to calculate values associated with these characteristics (for example, curve radius 1308, degree of curvature, curvature change rate, curve arc length, etc.) based on mathematical models and/or formulas. . In some embodiments, the processing unit 110 may execute the characteristic observation module 1404 to construct a mathematical model that expresses the curvature of the curve as a function of distance (for example, distance traveled). Such models may be subject to constraints, such as maximum lateral acceleration constraints (for example, in the range of 1.5 m/sec2 to 3.0 m/sec2) and/or maximum lateral acceleration derived constraints (for example, at 0.8 m/sec2 To 1.2 m/s2). These constraints can result in the maximum speed of the vehicle 200 (as a function of distance traveled).

[0278]

In step 1560, the processing unit 110 may execute the speed module 1406 to calculate an updated target speed based on the characteristics determined in step 1550.

The updated target speed may reflect the updated speed and/or direction of the vehicle 200 to safely perform any combination of the following: (i) entering the upcoming curve; (ii) navigating through the curve; or (iii) leaving the curve. The updated target speed can be determined using the technique described above in conjunction with FIG. 14 and step 1530. For example, the speed module 1406 may determine the updated target speed based on characteristics associated with the curve and based on mathematical models and/or formulas.

[0279]

In step 1570, the processing unit 110 may execute the speed module 1406 to adjust the speed of the vehicle 200 to the updated target speed determined in step 1560.

This adjustment can be accomplished using the technique described above in connection with step 1540. For example, based on the current speed of the vehicle 200 and the updated target speed determined in step 1560, the processing unit 110 may send electronic signals to the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200 (for example, via CAN Bus) to change the speed of the vehicle 200 to the updated target speed. At this point, the current speed of the vehicle 200 is not necessarily the same as the initial target speed determined in step 1530. For example, the processing unit 110 may have determined the updated target speed in step 1560, and at the same time adjust the speed of the vehicle 200 to the initial target speed. This may really exist when the difference between the speed of the vehicle 200 and the initial target speed is large (for example, greater than 10 km/hr) and the associated acceleration is small.

[0280]

In some embodiments, the system 100 may routinely update to the target speed based on continued observation of the characteristics of the curve (for example, based on image data of the portion of the vehicle 200 that depicts the curve).

In addition, after the system 100 navigates the vehicle 200 through the curve, the system 100 may cause the vehicle 200 to accelerate to a new target speed, which is suitable for the section of the road without curve, which is determined by the system through analysis. 100 is determined based on the map data, location data and/or image data acquired.

IPR2025-00943
Tesla EX1002 Page 1220

[0281]

Imitate the previous car

[0282]

The system 100 may provide driver assistance functionality for causing the vehicle 200 to imitate (or refuse to imitate) the preceding vehicle in different scenarios (such as when the current vehicle switches lanes, accelerates, or turns).

For example, the system 100 may be configured to detect the preceding vehicle by analyzing multiple images and determine one or more actions taken by the preceding vehicle.

In some scenarios, such as when the current vehicle turns at an intersection, the system 100 may be configured to cause the vehicle 200 to refuse to imitate the turn. In other scenarios, such as when the current vehicle turns at an intersection and does not change its course within the intersection, the system 100 may be configured to cause the vehicle 200 to imitate the turn. The system 100 may also be configured to cause the vehicle 200 to imitate one or more actions of the preceding vehicle based on the navigation history of the preceding vehicle.

[0283]

16A and 16B are schematic representations of the host vehicle 200a that imitates one or more actions of the preceding vehicle 200b on the road 1600a consistent with the disclosed embodiment.

As illustrated in FIGS. 16A and 16B, the host vehicle 200a may track the preceding vehicle 200b while traveling in the same lane of the road 1600a. The lane may have a left edge 1620, a right edge 1630, and a midpoint 1610, and w1 and w2 indicate the left half and right half of the lane, respectively. The host vehicle 200a and the preceding vehicle 200b may be positioned in the lane so that the vehicles are not centered on the midpoint 1610. As shown in FIG. 16A, for example, the host vehicle 200a may be centered to the right of the midpoint 1610 such that the distance c1 between the host vehicle 200a and the left edge 1620 is greater than the distance c2 between the host vehicle 200a and the right edge 1630. In contrast, the leading vehicle 200b can be centered on the left of the midpoint 1610, so that the distance c3 between the leading vehicle 200b and the left edge 1620 is smaller than the distance c4 between the leading vehicle 200b and the right edge 1630. The positions of the host vehicle 200a and the preceding vehicle 200b may have different configurations. For example, the host vehicle 200a and the preceding vehicle 200b may be positioned anywhere on the road 1600a. Further, although FIG. 16A depicts the front vehicle 200b having the camera system, the disclosed embodiments are not limited to the configuration in which the front vehicle 200b includes the camera system.

[0284]

The system 100 of the host vehicle 200a may cause the host vehicle 200a to imitate one or more actions of the preceding vehicle 200b.

As shown in FIG. 16B, for example, the host vehicle 200a may imitate the lane shift performed by the preceding vehicle 200b. After imitating the lane shift, the distance c1 may be equal to the distance c3, and the distance c2 may be equal to the distance c4. In some embodiments, the host vehicle 200a may mimic the lane shift performed by the preceding vehicle 200b, so that the distances c1 and c3 are not equal to the distances c2 and c4. For example, in the case where the host vehicle 200a imitates the left lane shift of the host vehicle 200b and (after the lane shift) is positioned to the left of the host vehicle 200a, c1 will be less than c3 and c2 will be greater than c4. Alternatively, in the case where the host vehicle 200a imitates the left lane displacement of the host vehicle 200b and is positioned to the right of the host vehicle 200a, c1 will be greater than c3 and c2 will be less than c4.

[0285]

16C and 16D are schematic representations of the host vehicle 200a that imitates one or more actions of the preceding vehicle 200b on the road 1600b consistent with the disclosed embodiment.

As illustrated in FIG. 16C, the host vehicle 200a may travel behind the preceding vehicle 200b, and the preceding vehicle 200b may turn left at the intersection. The system 100 of the host vehicle 200a may cause the host vehicle 200a to imitate a left turn of the preceding vehicle 200b, as shown in FIG. 16D. The system 100 may cause the host vehicle 200a to change lanes while

turning in the intersection, navigation history associated with the preceding vehicle 200b, the relative speed difference between the host vehicle 200a and the preceding vehicle 200b, etc., based on various factors. The vehicle 200a imitates a left turn. In addition, the host vehicle 200a may mimic the preceding vehicle 200b under various conditions, such as based on whether the preceding vehicle 200b travels along the same route or the same part of the route along which the host vehicle 200a travels. Further details on the scenario in which the system 100 can cause the host vehicle 200a to imitate one or more actions of the preceding vehicle 200b are described below in conjunction with FIGS. 17-22.

[0286]

In some embodiments, the system 100 may be configured to determine whether to mimic the specific movement of the preceding vehicle.

For example, the system 100 may determine that certain movements of the preceding vehicle do not need to be imitated because such movements do not affect the route of the preceding vehicle. For example, the system 100 may determine that small changes in movement (such as movement in a specific lane) need not be imitated. In other embodiments, the system 100 may implement a smooth operation to filter out a small amount of movement of the vehicle in front. Smooth operation may cause the host vehicle to perform more significant and/or important movements of the preceding vehicle (for example, changing lanes), while filtering out a smaller amount of movement that does not affect the overall route of the preceding vehicle.

[0287]

FIG. 17 is an exemplary block diagram of a memory consistent with the disclosed embodiments, the memory being configured to store instructions for performing one or more operations.

As shown in FIG. 17, the memory 140 can store a lane restriction module 1702, an action detection module 1704, and an action response module 1706. The disclosed embodiments are not limited to the memory 140 of any particular configuration. Further, the processing unit 110 can execute instructions stored in any of the modules 1702-1706 (included in the memory 140).

[0288]

In one embodiment, the lane restriction module 1702 may store software instructions (such as computer vision software) that, when executed by the processing unit 110, determine one or more restrictions associated with the lane of travel of the host vehicle 200a.

The lane constraints may include a midpoint 1610, a left edge 1620, and a right edge 1630 of the lane. The lane constraints may also include one or more lines or other symbols marked on the surface of the road 1600a or 1600b. The processing unit 110 may execute the lane restriction module 1702 to determine the restriction by, for example, analyzing one or more sets of images and/or using map and/or location data (e.g., data stored in the map database 160). The image group may be acquired by one or more of the image capture devices 122, 124, and 126. Based on this analysis, the system 100 (eg, via the processing unit 110) may cause the host vehicle 200a to travel within the lane defined by the left edge 1620 and the right edge 1630. In some embodiments, if one of the distances c1 or c2 is less than a predetermined threshold, the system 100 may cause the host vehicle 200a to move away from the left edge 1620 or the right edge 1630.

[0289]

In one embodiment, the action detection module 1704 may store software instructions (such as computer vision software) that, when executed by the processing unit 110, detect one or more actions taken by the preceding vehicle 200b traveling in front of the host vehicle 200a .

The preceding vehicle 200b may travel in the same or different lane as the host vehicle 200a. The processing unit 110 may execute the action detection module 1704 to detect one or more actions taken by the preceding vehicle 200b by, for example, analyzing one or more sets of images acquired by one or more of the image capturing devices 122, 124, and 126 . Based on this analysis, for example, the processing unit 110 may determine that the preceding vehicle 200b has changed lanes, has turned, accelerated, decelerated, applied brakes, and the like. As another example, the preceding vehicle 200b may perform a maneuver such that the first offset distance of the preceding vehicle 200b on the side adjacent to the first lane constraint is different from the second offset distance of the preceding vehicle 200b on the side adjacent to the second lane constraint. Offset distance. For example, the processing unit 110 may perform such analysis based on the technique described above in conjunction with FIGS. 5A to 5D.

IPR2025-00943
Tesla EX1002 Page 1222

[0290]

In one embodiment, the action response module 1706 may store software instructions (such as computer vision software) that, when executed by the processing unit 110, are based on the preceding vehicle 200b detected by the system 100 (for example, via the execution of the action detection module 1704). The one or more actions taken determine the one or more actions taken by the host vehicle 200a.

For example, the processing unit 110 may execute the action response module 1706 to determine whether to imitate one or more actions of the preceding vehicle 200b. The determination may be based on the nature of the actions of the preceding vehicle 200b (for example, turning, lane shifting, etc.), information associated with the host vehicle 200a (for example, the speed from the side of the road, distances c1 and c2, etc.), Road and environmental conditions (for example, potholes, heavy rain or wind, etc.), etc. In the case that the processing unit 110 determines to imitate one or more actions of the preceding vehicle 200b, the processing unit 110 may accomplish this by the following means: for example, to the throttle system 220, the braking system 230 and/or the steering system 240 of the host vehicle 200a. Electronic signals are sent (e.g., via a CAN bus) to trigger turns, lane shifts, speed changes and/or direction changes of host vehicle 200a. The processing unit 110 may use the techniques described above in conjunction with FIGS. 4 to 7 to cause the host vehicle 200a to imitate one or more actions. In order to imitate one or more actions, the system 100 may provide control signals to one or more of the throttle system 220, the braking system 230, and the steering system 240 (for example, by causing acceleration, turning, lane shift, etc.) Etc.) Navigation vehicle 200. Further, one or more actuators may control the throttle system 220, the braking system 230, and/or the steering system 240.

[0291]

FIG. 18 is a flowchart showing an exemplary process for causing the host vehicle 200a to imitate one or more actions of the preceding vehicle 200b consistent with the disclosed embodiment.

In step 1810, the processing unit 110 may receive a plurality of images via the data interface 128 between the processing unit 110 and the image acquisition unit 120. For example, a camera (such as the image capturing device 122) included in the image acquisition unit 120 may capture a plurality of images and send the images to the processing unit 110 via a digital connection (eg, USB, wireless, Bluetooth, etc.). In some embodiments, the processing unit 110 may receive more than one multiple images via multiple data interfaces. For example, the processing unit 110 may receive multiple images from each of the image capture devices 122, 124, 126, and each of the image capture devices 122, 124, 126 may have an associated data interface for communicating with the processing unit 110 data. The disclosed embodiments are not limited to any specific data interface structure or protocol.

[0292]

In steps 1820 and 1830, the processing unit 110 may execute the lane restriction module 1702 to determine constraints associated with the travel lane of the host vehicle 200a on each side of the host vehicle 200a.

The processing unit 110 may determine the constraint by analyzing the image received in step 1810. Thus, in step 1820, the processing unit 110 may determine the first lane constraint on the first side of the host vehicle 200a (for example, the distance between the left side of the host vehicle 200a and the left edge 1620 of the lane, c1), and In step 1830, the second lane constraint on the second side opposite to the first side is determined (e.g., the distance between the right side of the host vehicle 200a and the right edge 1620 of the lane, c2). In some embodiments, the processing unit 110 may use a map and/or position data indicating the position of the host vehicle 200a relative to the midpoint 1610, left edge 1620, and right edge 1630 of the lane (for example, stored in the map database 160). Data) to determine constraints. The processing unit 110 may determine the constraint based on analyzing both the image received in step 1810 and the map and/or location data; doing so may increase the confidence level associated with the constraint.

[0293]

In step 1840, the processing unit 110 may execute the lane restriction module 1702 to cause the host vehicle 200a to travel within the lane restriction determined in steps 1820 and 1830.

For example, the left side of the host vehicle 200a may be positioned to the left of the left edge 1620, or the right side of the host vehicle 200a may be positioned to the right of the right edge 1630. In such a scenario, the processing unit 110 may send

electronic signals (for example, via the CAN bus) to the throttle system 220, the braking system 230, and/or the steering system 240 of the host vehicle 200a to cause the host vehicle 200a to adjust its position and Travel within the lane constraints.

[0294]

In step 1850, the processing unit 110 may execute the motion detection module 1704 to locate the preceding vehicle 200b within the image received in step 1810 by analyzing the image.

The processing unit 110 may use the technique described above in conjunction with FIG. 5A and FIG. 5B to detect the presence of the vehicle in the image. For each detected vehicle, the processing unit 110 may construct a multi-frame model of the position, speed (for example, speed and direction), and acceleration of the detected vehicle relative to the host vehicle 200a. The processing unit 110 can construct such a model using the technique described above in conjunction with FIG. 5A and FIG. 5B. In addition, the processing unit 110 may determine the lane of travel associated with each detected vehicle, based on, for example, lanes or other road markings and the lane constraints determined in steps 1820 and 1830. The processing unit 110 may determine that the preceding vehicle 200b is the closest detected vehicle and is traveling in the same lane as the host vehicle 200a. In step 1860, the processing unit 110 can execute the action detection module 1704 to determine one or more actions taken by the preceding vehicle 200b in the image detected in step 1850. For example, the processing unit 110 may determine: the preceding vehicle 200b has changed lanes, has turned, accelerated, decelerated, and so on. In addition, the processing unit 110 may analyze the image to determine characteristics associated with the preceding vehicle 200b, such as its (e.g., relative to the host vehicle 200a) speed and position. These determinations can be based on the techniques described above in conjunction with FIGS. 5A and 5B. In some embodiments, the processing unit 110 may analyze the image to determine characteristics associated with the action performed by the preceding vehicle 200b. For example, in a situation where the preceding vehicle 200b is turning, the processing unit 110 may determine the position information of the preceding vehicle 200b just before starting the turn, the position information of the preceding vehicle 200b when the turn is completed, the turning radius, and the description of the position of the preceding vehicle 200b when the turn is completed. The velocity profile of the velocity change (for example, over a period of time).

[0295]

In step 1870, the processing unit 110 may execute the action response module 1706 to cause the host vehicle 200a to imitate one or more actions of the preceding vehicle 200b determined in step 1850.

For example, the processing unit 110 may send electronic signals (for example, via the CAN bus) to the throttle system 220, the braking system 230, and/or the steering system 240 of the host vehicle 200a to trigger the turning, lane shift, and speed of the host vehicle 200a. Change and/or change of direction. In addition, the processing unit 110 may construct a set of waypoints to guide the position and speed of the host vehicle 200a to change. The waypoints may be based on lane markings (e.g., midpoint 1610, left edge 1620, right edge 1630, etc.) represented by a polynomial (e.g., third polynomial) and target coordinates associated with the target vehicle (such as the preceding vehicle 200b) Wait). The processing unit 110 may adjust the coefficients of the polynomial so that the path points exceed the target coordinates. The processing unit 110 may also offset the waypoint (for example, with respect to the left edge 1620 and the right edge 1630) based on a predetermined offset amount. When the processing unit 110 adjusts the polynomial coefficients and offsets the path points, the processing unit 110 may also calculate local curvature information for each obtained section formed between the path points.

[0296]

In some embodiments, the processing unit 110 (eg, via the execution action response module 1706) may cause the host vehicle 200a to mimic the preceding vehicle 200b within a certain position, speed, and/or acceleration constraint associated with the two vehicles.

For example, the constraint may require a predetermined minimum distance (eg, a safety distance) between the host vehicle 200a and the preceding vehicle 200b or a predetermined maximum difference in speed and/or acceleration between the vehicles. Thus, if the action performed by imitating the preceding vehicle 200b would violate one or more constraints (for example, imitating acceleration would cause the host vehicle 200a to be too close to the preceding vehicle 200b), the processing unit 110 may cause the host vehicle 200a to refuse to imitate the preceding vehicle 200b. One or more actions. In addition, the processing unit 110 may cause the host vehicle 200a to follow the speed profile of the preceding vehicle 200b and maintain a predetermined distance or a safe distance from the preceding vehicle 200b even while mimicking the action performed by the preceding vehicle 200b (for example, lane shift). The processing unit 110 can accomplish this by causing the

host vehicle 200a to: (i) continuously shorten the distance to the preceding vehicle 200b when the distance is greater than the predetermined minimum distance; (ii) continuously reduce the speed difference between the host vehicle 200a and the preceding vehicle 200b Gap; and (iii) matching the acceleration of the preceding vehicle 200b to a degree consistent with predetermined constraints. In order to control to what extent these actions should be performed, the processing unit 110 may associate each action with a weight in real time based on the current predicted distance from the preceding vehicle 200b. Thus, for example, when the distance to the preceding vehicle 200b decreases, a greater weight may be added to the deceleration instead of matching the acceleration of the preceding vehicle 200b.

[0297]

As another example, in a case where the host vehicle 200a approaches the preceding vehicle 200b at a speed greater than the speed of the preceding vehicle 200b, the processing unit 110 may cause the host vehicle 200a to gradually decelerate, so that when the distance between the two vehicles is equal to the predetermined safety distance ((Such as the minimum following distance), the speed of the two vehicles is the same.

In another example, when the preceding vehicle 200b suddenly appears in front of the host vehicle 200a, causing the distance between the two vehicles to be less than the minimum following distance, the processing unit 110 may cause the host vehicle 200a to gradually decelerate to avoid getting closer to the preceding vehicle. 200b, at the same time increase the distance between vehicles until the distance reaches the minimum following distance.

[0298]

19 is a flowchart showing an exemplary process 1900 for causing host vehicle 200a to refuse to imitate a turn of the preceding vehicle 200b, consistent with the disclosed embodiment.

In step 1910, as described above in conjunction with step 1810 of FIG. 18, the processing unit 110 may receive multiple images via the data interface 128 between the processing unit 110 and the image acquisition unit 120. In step 1920, as described above in conjunction with step 1850 of FIG. 18, the processing unit 110 may execute the motion detection module 1704 to locate the preceding vehicle 200b in the image received in step 1910 by analyzing the image. In step 1925, the processing unit 110 may execute the action detection module 1704 to use the technique described above in conjunction with step 1860 in FIG. 18 to determine whether the vehicle 200b in front is turning at the intersection. If it is determined that the preceding vehicle 200b is turning at the intersection (step 1925, yes), then in step 1930, the processing unit 110 may cause the host vehicle 200a to refuse to imitate the turn of the preceding vehicle 200b. Otherwise (step 1925, No), the process 1900 ends.

[0299]

20 is a flowchart showing an exemplary process 2000 for causing the host vehicle 200a to imitate or reject the turn of the preceding vehicle 200b consistent with the disclosed embodiment.

In step 2010, as described above in conjunction with step 1810 of FIG. 18, the processing unit 110 may receive multiple images via the data interface 128 between the processing unit 110 and the image acquisition unit 120. In step 2020, as described above in conjunction with step 1850 in FIG. 18, the processing unit 110 may execute the motion detection module 1704 to locate the vehicle 200b in front in the image received in step 2010 by analyzing the image. In step 2025, the processing unit 110 may execute the action detection module 1704 to use the technique described above in conjunction with step 1860 in FIG. 18 to determine whether the preceding vehicle 200b is turning at the intersection. If it is determined that the preceding vehicle 200b is turning at the intersection (step 2025, YES), then in step 2035, the processing unit 110 may execute the motion detection module 1704 to determine whether the preceding vehicle 200b is at step 1860 using the technique described above in conjunction with step 1860 of FIG. Change lanes within the intersection. Otherwise (step 2025, No), the process 2000 ends. If it is determined that the preceding vehicle 200b is changing lanes within the intersection (step 2035, YES), then in step 2050, the processing unit 110 may cause the host vehicle 200a to refuse to imitate the turn of the preceding vehicle 200b. Otherwise (step 2035, No), in step 2040, the processing unit 110 may cause the host vehicle 200a to imitate the turn of the preceding vehicle 200b.

[0300]

21 is a flowchart showing another exemplary process 2100 for causing the host vehicle 200a to imitate one or more actions of the first preceding vehicle consistent with the disclosed embodiment.

In step 2110, as described above in conjunction with step 1810 in FIG. 18, the processing unit 110 may receive multiple images via the data interface 128 between the processing unit 110 and the image acquisition unit 120. In steps 2120 and 2130, as described above in conjunction with step 1850 of FIG. 18, the processing unit 110 may execute the motion detection module 1704 to locate the first and second vehicles ahead in the image received in step 2110 by analyzing the image. In step 2135, the processing unit 110 can execute the motion detection module 1704 to determine whether the turning radius of the first preceding vehicle path is smaller than the second preceding vehicle path using the technique described above in conjunction with step 1860 in FIG. 18. If it is determined that the turning radius of the first preceding vehicle path is smaller than the second preceding vehicle path (step 2135, YES), then in step 2140, the processing unit 110 may cause the host vehicle 200a to imitate the turn of the first preceding vehicle. Otherwise (step 2135, No), the process 2100 ends.

[0301]

22 is a flowchart showing an exemplary process 2200 for causing the host vehicle 200a to imitate one or more actions of the preceding vehicle 200b based on the navigation history, consistent with the disclosed embodiment.

In step 2210, as described above in conjunction with step 1810 of FIG. 18, the processing unit 110 may receive multiple images via the data interface 128 located between the processing unit 110 and the image acquisition unit 120. In step 2220, as described above in conjunction with step 1850 in FIG. 18, the processing unit 110 may execute the motion detection module 1704 to locate the vehicle 200b in front within the image received in step 2210 by analyzing the image.

[0302]

In step 2230, the processing unit 110 may determine the position information of the preceding vehicle 200b by analyzing the image received in step 2210.

The position information of the preceding vehicle 200b may indicate the position of the preceding vehicle 200b relative to the host vehicle 200a. The processing unit 110 may analyze the image to obtain the position information of the preceding vehicle 200b and one or more actions taken by the preceding vehicle 200b by using the techniques described above in conjunction with FIGS. 5A to 5D.

[0303]

In step 2240, the processing unit 110 may compare the position information of the preceding vehicle 200b determined in step 2230 with predetermined map data (for example, data stored in the map database 160).

For example, by comparing the location information with predetermined map data, the processing unit 110 can obtain the location coordinates of the preceding vehicle 200b. In addition, the processing unit 110 may use predetermined map data as an indicator of the confidence level associated with the location information determined in step 2230.

[0304]

In step 2250, the processing unit 110 may create a navigation history associated with the preceding vehicle 200b by, for example, tracking actions taken by the preceding vehicle 200b.

The processing unit 110 may associate the action taken by the preceding vehicle 200b (for example, determined according to the analysis image in step 2230) with the predetermined map data in step 2240. Thus, the navigation history may indicate that the preceding vehicle 200b performing certain actions (for example, turning, lane shifting, etc.) is at a certain position coordinate.

[0305]

In step 2260, the processing unit 110 may execute the action response module 1706 to determine whether to cause the host vehicle 200a to imitate one or more of the preceding vehicle 200b based on the navigation history created in step 2250 and using the technique described above in conjunction with step 1870 of FIG. 18 Actions.

For example, even in the case where no turn is detected according to the analysis image (for example, in step 2230), when the navigation history indicates that the preceding vehicle 200b has successfully navigated at least one previous turn, the processing unit 110 (for example, via the action response module 1706) The host vehicle 200a may be caused to imitate the

Tesla EX1002 Page 1226 IPR2025-00943

IPR2025-00943 Tesla EX1002 Page 1227

turn of the preceding vehicle 200b. More generally, in a case where the navigation history indicates that the preceding vehicle 200b has successfully performed such an action, the processing unit 110 may cause the host vehicle 200a to imitate the action taken by the preceding vehicle 200b.

[0306]

Navigate the vehicle to overtake another vehicle

[0307]

The system 100 may provide driver assistance functionality for monitoring the vicinity of the vehicle 200 and aborting the lane change if the lane change is deemed to cause a risk such as a collision with another vehicle (eg, a target vehicle).

For example, if it is determined that the target vehicle is located in a different lane from the vehicle 200, the system 100 may enable the vehicle 200 to complete the overtaking of the target vehicle.

If, before completing the overtaking, the system 100 determines that the target vehicle is entering the driving lane of the vehicle 200, the system 100 may cause the vehicle 200 to abort the overtaking. For example, aborting an overtaking may include causing the vehicle 200 to brake or causing the vehicle 200 to stop accelerating.

[0308]

As illustrated in FIG. 23, the vehicle 200 can travel on the road 2300.

The vehicle 200 may be equipped with the system 100 and may implement any of the disclosed embodiments to recognize lane constraints and operate the vehicle 200 within the lane constraints. The system 100 may also facilitate lane changes made by the vehicle 200. For example, the system 100 may determine whether the situation allows the vehicle 200 to change lanes safely, and if the situation changes during the lane change, cause the vehicle 200 to abort the lane change.

[0309]

As shown in FIG. 23 and as previously discussed with respect to FIGS. 8A to 8D, the vehicle 200 may include a first vehicle side 802, and the first vehicle side 802 may be a first distance 805 from the first lane constraint 2302.

Similarly, the vehicle 200 may include a second vehicle side 812 opposite the first vehicle side 802, and the second vehicle side 812 may be a second distance 815 from the second lane constraint 2304. In this way, the first lane constraint 810 and the second lane constraint 2304 may define the lane 2330 in which the vehicle 200 may travel.

[0310]

The system 100 may determine the first lane constraint 2302 and the second lane constraint 2304 based on multiple images acquired by one or more of the image capture devices 122-126.

According to some embodiments, the first lane constraint 2302 and/or the second lane constraint 2304 may be identified by visible lane boundaries (including lines marked on the road surface). Additionally or alternatively, the first lane constraint 2302 and/or the second lane constraint 2304 may include the curb of the road surface. According to some embodiments, the system 100 can determine the first lane constraint 2302 and/or the second lane constraint 2304 by identifying the midpoint of the road width 2320 (or any other suitable road feature that the vehicle 200 can use as a navigation reference). For example, the system 100 can identify lane constraints in this way when the lines specifying the lanes of the road are not painted or otherwise labeled.

[0311]

When overtaking the target vehicle 2310, the lane change of the vehicle 200 may be complicated due to the change of the trajectory of the target vehicle 2330.

Thus, the system 100 can identify the target vehicle 2310 based on analysis of multiple images acquired by one or more of the image capture devices 122-126. The system 100 may also identify the second lane 2340 in which the target vehicle 2310 travels. Further, the system 100 may determine whether the second lane 2340 is different from the lane 2330 in which the

vehicle 200 is traveling. If the second lane 2340 is different from the lane 2330, the system 100 may be configured to enable the vehicle 200 to pass the target vehicle 2310.

[0312]

The system 100 may also monitor the target vehicle 2310 based on multiple images.

For example, the system 100 may monitor the position of the target vehicle 2310, such as monitoring the position of the target vehicle 2310 relative to the vehicle 200, relative to the first lane constraint 2302 and/or the second lane constraint 2304, absolute position, or relative to another reference point. Location. Additionally or alternatively, monitoring the location of the target vehicle 2310 may include estimating the speed of the target vehicle 2310.

[0313]

If the system 100 determines that the target vehicle 2310 is entering the lane 2330 in which the vehicle 200 is traveling, the system 100 may cause the vehicle 200 to abort the overtaking of the target vehicle 2310.

For example, the system 100 may cause the vehicle 200 to stop accelerating or apply braking. Additionally or alternatively, the system 100 may issue a voice notification, such as informing the driver of the vehicle 200 to abort the lane change.

[0314]

FIG. 24A is an exemplary block diagram of the memory 140 consistent with the disclosed embodiments, and the memory 140 may store instructions for performing one or more operations.

As illustrated in FIG. 24A, the memory 140 may store one or more modules for identifying the target vehicle and the responses described herein. For example, the memory 140 may store the lane restriction module 2400, the target vehicle acquisition module 2410, and the action module 2415.

[0315]

In one embodiment, the lane restriction module 2400 may store instructions, and when executed by the processing unit 110, the instructions may detect and define the first lane restriction 2302 and the second lane restriction 2304.

For example, the processing unit 110 may execute the lane offset module 2400 to process a plurality of images received from at least one image capturing device 122-124 and detect the first lane constraint 2302 and the second lane constraint 2304. As discussed above, this may include identifying painted lane lines or markers and/or measuring the midpoint of the road surface. The processing unit 110 may perform this analysis based on the technique described above in conjunction with FIGS. 5A to 5D.

[0316]

In some embodiments, the target acquisition module 2410 may store instructions, and when executed by the processing unit 110, the instructions may detect the existence of the target vehicle 2310 and monitor the location of the target vehicle 2310.

For example, the target acquisition module 2410 may process multiple images to detect and monitor the target vehicle 2310. Additionally or alternatively, when the target acquisition module 2410 is executed, the processing unit 110 may receive information from another module or other system that indicates the existence and/or location of the target vehicle 2310. The processing unit 110 may perform this analysis based on the technique described above in conjunction with FIGS. 5A to 5D.

[0317]

The processing unit 110 may be configured to determine the existence of one or more objects (such as a target vehicle 2310 in the vicinity of the vehicle 200).

For example, the target vehicle 2310 may include another vehicle traveling near the vehicle 200, such as a car, a truck, or a motorcycle. For each target vehicle 2310, the processing unit 110 may determine an offset profile. The offset profile may include determining whether the target vehicle 2310 will be located within a predetermined range of the vehicle 200 in its current predicted position. If the processing unit 110 determines that the target vehicle 2310 is located or will be located within the predetermined range of the vehicle 200, the processing unit may determine whether there is any in the first lane constraint 810

and the second lane constraint 820 in order for the vehicle 200 to bypass the target vehicle 2310. enough space. If there is not enough space, the processing unit 110 may perform a lane change. If there is enough space, the processing unit 110 may determine whether there is enough time to bypass the target vehicle 2310 before the distance between the vehicle 200 and the target vehicle 2310 is shortened. If there is enough time, the processing unit 110 may initiate or activate an offset maneuver to bypass the target vehicle 2310. The offset maneuver may be performed by the processing unit 110 so that the movement of the vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 m/sec to 0.75 m/sec. The offset maneuver may be performed such that the maximum amplitude of the offset maneuver will occur when the distance between the vehicle 200 and the object 825 is shortened. If there is not enough time, the processing unit 110 may initiate braking and then initiate an offset maneuver.

[0318]

Further, the processing unit 110 may execute the action module 2415 to suspend the plan of the vehicle 200 to overtake the target vehicle 2310.

If the target vehicle 2310 is not in the lane 2330 of the vehicle 200, the instructions included in the action module 1230 may enable the vehicle 200 to pass the target vehicle 2330. However, if the execution target vehicle acquisition module 2410 indicates that the target vehicle 2310 is changing lanes and entering the lane 2330 of the vehicle 200, the processing unit 110 may execute the instruction included in the action module 2415 to stop the lane change of the vehicle 200. For example, the processing unit 110 may send an electronic signal to the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200 to stop the lane change. Additionally or alternatively, the processing unit may cause a voice notification to be issued.

[0319]

FIG. 24B illustrates a flowchart of an exemplary process 2420 for navigating a vehicle in an intruding vehicle consistent with the disclosed embodiment.

The process 2420 may identify the lane constraints that define the travel lane in which the vehicle (for example, vehicle 200) travels, identify and monitor the target vehicle (for example, vehicle 2310), and enable the vehicle to pass the target vehicle or if it is determined that the target vehicle is entering the travel of the vehicle The lane is suspended for overtaking.

[0320]

In step 2430, at least one of the image capturing devices 122, 124, and/or 126 may acquire multiple images of the area near the vehicle 200.

For example, a camera (such as the image capturing device 122) included in the image acquisition unit 120 may capture a plurality of images and send the images to the processing unit 110 via a digital connection (eg, USB, wireless, Bluetooth, etc.). In some embodiments, the processing unit 110 may receive more than one multiple images via multiple data interfaces. For example, the processing unit 110 may receive multiple images from each of the image capture devices 122, 124, 126, and each of the image capture devices 122, 124, 126 may have an associated data interface for communicating with the processing unit 110 data. The disclosed embodiments are not limited to any specific data interface structure or protocol.

[0321]

Next, in step 2440, the processing unit 110 may use the multiple images received via the data interface 128 to determine the first lane constraint on the first vehicle side 802 and the second lane constraint on the second vehicle side 812.

For example, as part of implementing step 2440, the processing unit 110 may execute instructions of the lane restriction module 910. Further, as a part of determining the first lane constraint and the second lane constraint, the processing unit 110 may use one or more of the processes discussed above in conjunction with FIGS. 5A to 5D.

[0322]

In step 2450, the processing unit 110 may obtain the target vehicle 2310.

Tesla EX1002 Page 1229 IPR2025-00943

In order to obtain the target vehicle 2310, the processing unit 110 may execute the target vehicle obtaining module 2410. For example, the processing unit 110 may execute instructions to analyze the acquired image to identify the second lane 2340 traveled by the target vehicle 2310 and determine whether the target vehicle 2310 travels in the lane 2330 traveled by the vehicle 200. This determination may be based on the technique described above in conjunction with FIGS. 5A to 5D.

[0323]

In step 2455, the processing unit 110 may determine whether the second lane 2340 traveled by the target vehicle 2310 is different from the lane 2330 traveled by the vehicle 200.

[0324]

In step 2460, if the second lane 2340 is different from the lane 2330 in which the vehicle 200 travels, the processing unit 110 may enable the vehicle 200 to pass the target vehicle 2310.

In order to enable the vehicle 200 to overtake the target vehicle 2310, the processing unit may execute the action module 2415.

[0325]

Next, in step 2465, the processing unit 110 may execute the target acquisition module 2410 to monitor the target vehicle 2310.

This may include estimating the speed of the target vehicle 2310 and determining the location of the target vehicle 2310, such as relative to the vehicle 200.

[0326]

Step 2470 may include determining whether the target vehicle 2310 has entered the lane 2330.

If the processing unit 110 determines that the target vehicle 2310 has not entered the lane 2330, then in step 2480, the processing unit 110 may allow the vehicle 200 to complete the overtaking.

[0327]

In step 2490, if the processing unit 110 determines that the target vehicle 2310 has entered the lane 2330 (or is entering the lane 2330 or is otherwise on a trajectory that will bring the vehicle 2310 into the lane 2330), the processing unit 110 may execute the action module 2415 To stop the transcendence.

For example, this may include issuing a voice notification to the driver of the vehicle 200 via the speaker 360.

Additionally or alternatively, step 2490 may include executing instructions included in the action module 2415 to cause the processing unit 110 to send electronic signals to the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200, thereby suspending Lane change.

[0328]

Navigate the vehicle to avoid intrusion into traffic

[0329]

The system 100 may provide driver assistance functionality for monitoring the vicinity of the vehicle 200 and responding to the presence of another vehicle invading the vehicle 200.

Monitoring the vicinity of the vehicle 200 may include monitoring a collision threshold.

For example, the collision threshold may include the time to collision or the minimum predetermined distance between the vehicle 200 and other traffic. The system 100 may further cause the vehicle 200 to take an evasive action to avoid intruding into traffic. For example, the avoidance action may include braking, changing lanes, or otherwise changing the course of the vehicle 200.

[0330]

As illustrated in FIG. 25, the vehicle 200 may travel on a road 2500.

The vehicle 200 may include a system 100 for detecting and responding to intrusive traffic. For example, the vehicle 200 may determine that another vehicle is entering the lane of travel of the vehicle 200, or otherwise determine that the other vehicle is crossing the collision threshold. The vehicle 200 may react to the intruding vehicle by, for example, changing the route, braking, or accelerating.

[0331]

As discussed with respect to FIGS. 8A to 8D and as shown in FIG. 25, the vehicle 200 may include a first vehicle side 802, and the first vehicle side 802 may be a first distance 805 from the first lane constraint 2502.

Similarly, the vehicle 200 may include a second vehicle side 812 opposite the first vehicle side 802, and the second vehicle side 812 may be a second distance 815 from the second lane constraint 2504. In this way, the first lane constraint 2502 and the second lane constraint 2504 may define the lane 2505 in which the vehicle 200 may travel.

[0332]

The system 100 may determine the first lane constraint 2502 and the second lane constraint 2504 based on multiple images acquired by one or more of the image capture devices 122-126.

According to some embodiments, the first lane constraint 2502 and/or the second lane constraint 2504 may be identified by visible lane boundaries, such as lines marked on the road surface. Additionally or alternatively, the first lane constraint 2502 and/or the second lane constraint 2504 may include the curb of the road surface. According to some embodiments, the system 100 may determine the first lane constraint 2502 and/or the second lane constraint 2502 by identifying the midpoint 2520 of the road width (or any other suitable road feature vehicle 200 may be used as a navigation reference). The system 100 can identify lane constraints in this way when, for example, the line specifying the lane of the road has not been painted or otherwise labeled.

[0333]

Detecting the first lane constraint 2502 and/or the second lane constraint 2504 may include the processing unit 110 determining their three-dimensional models in the camera coordinate system.

For example, the three-dimensional model of the first lane constraint 2502 and/or the second lane constraint 2504 can be described by a third polynomial. In addition to the three-dimensional modeling of the first lane constraint 2502 and/or the second lane constraint 2504, the processing unit 110 may estimate motion parameters, such as the speed, yaw and pitch rate, and acceleration of the vehicle 200. Optionally, the processing unit may detect static and moving vehicles and their position, heading, speed and acceleration (all relative to vehicle 200). The processing unit 110 may determine a road elevation model to convert information acquired from a plurality of images into a three-dimensional space.

[0334]

What the driver of the vehicle 200 cannot assume is that if it is safe to shift the traffic into the lane 2505 in which the vehicle 200 is traveling, other traffic traveling in adjacent or nearby adjacent lanes will only change lanes.

Thus, it may be desirable to consider and monitor traffic traveling in a nearby lane to detect lateral intrusion into the lane 2505 of the vehicle 200 so that the vehicle 200 can react to the lateral intrusion to avoid a collision. The system 100 may identify the intruding vehicle 2510 based on the multiple images and determine that the intruding vehicle 2510 is approaching the vehicle 200. The system 100 can then generate actions to consider the intrusion into the vehicle. For example, the system 100 may cause the vehicle 200 to maintain the current speed and travel within the first lane constraint 2502 and the second lane constraint 2504 such that the first distance 2530 on the first vehicle side 802 where the intruding vehicle 2510 approaches is greater than the second distance 2540 .

[0335]

Tesla EX1002 Page 1231 IPR2025-00943

The processing unit 110 may be configured to determine the presence of one or more intruding vehicles 2510 in the vicinity of the vehicle 200.

For each intruding vehicle 2510, the processing unit 110 may determine an offset profile. The offset profile may include determining whether the invading vehicle 2510 will be located within a predetermined range of the vehicle 200 in its current predicted position. If the processing unit 110 determines that the intruding vehicle 2510 is located or will be located within the predetermined range of the vehicle 200, the processing unit may determine whether there are any within the first lane constraint 2502 and the second lane constraint 2504 in order for the vehicle 200 to bypass the intruding vehicle 2510 enough space. If there is not enough space, the processing unit 110 may perform a lane change. If there is enough space, the processing unit 110 may determine whether there is enough time to bypass the intruding vehicle 2510 before the distance between the vehicle 200 and the intruding vehicle 2510 is shortened.

[0336]

For example, if there is enough time, the processing unit 110 may initiate or activate an offset maneuver to bypass the invading vehicle 2510.

The offset maneuver may be performed by the processing unit 110 so that the movement of the vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 m/sec to 0.75 m/sec. The excursion maneuver may be performed such that the maximum amplitude of the excursion maneuver will occur when the distance between the vehicle 200 and the invading vehicle 2510 is shortened. If there is not enough time, the processing unit 110 may initiate braking (e.g., via sending an electronic signal to the braking system 230), and then initiate an offset maneuver (e.g., by sending an electronic signal to the throttle system 220 and/or the steering system 240). Signal).

[0337]

The system 100 may determine based on multiple images acquired by one or more of the image capture devices 122-126 that the intruding vehicle 2510 has crossed at least one collision threshold.

For example, the at least one collision threshold may include the minimum predetermined distance 2550 between the vehicle 200 and the invading vehicle 2510. For example, as illustrated in FIG. 25, the intruding vehicle 2510 is shown as exceeding the minimum predetermined distance 2550. Additionally or alternatively, the at least one collision threshold may include the time when the vehicle 200 and the invading vehicle 2510 are about to collide. The time-based collision threshold may be determined based on the following: the distance between the vehicle 200 and the intruding vehicle 2510, the speed of the vehicle 200 and/or the intruding vehicle 2510, the lateral speed of the vehicle 200 and/or the intruding vehicle 2510, and/or the vehicle 200 And/or the acceleration of the intruding vehicle 2510. The collision threshold based on the collision time can be set to allow the system 100 to have enough time to effectively perform the evasive action to avoid the collision.

[0338]

If the system 100 determines that at least one collision threshold has been crossed, the system 100 may cause the vehicle 200 to take an evasive action.

For example, this may include causing other subsystems of the vehicle 200 to operate. Thus, the system 100 may be operated to cause the speed of the vehicle 200 to increase by communicating with the throttle system 220, the braking system 230, and/or the steering system 240 to change the speed and/or direction of the vehicle 200.

[0339]

The evasive action may include causing the vehicle 200 to change lanes.

It may be desirable to respond in a situation where, for example, the vehicle 200 is traveling on a multi-lane highway, and the system 100 determines that changing lanes away from the invading vehicle 2510 will avoid a collision. The avoidance action may include causing the vehicle 200 to change its course. Additionally or alternatively, the avoidance action may include changing the route of the vehicle 200. For example, in order to take an evasive action, the processing unit 110 may send an electronic signal to the throttle system 220, the braking system 230, and/or the steering system 240.

[0340]

Tesla EX1002 Page 1232 IPR2025-00943

The system 100 may also be configured to cause the vehicle 200 to take an action after the dodging action has been completed.

For example, the system 100 may cause the vehicle 200 to return to the midpoint 2520 of the lane 2505 in which the vehicle 200 travels. Additionally or alternatively, the system 100 may cause the vehicle 200 to return to a preset speed after braking in response to the intruding vehicle 2510.

[0341]

FIG. 26 is an exemplary block diagram of a memory 140 consistent with the disclosed embodiment, and the memory 140 may store instructions for detecting and responding to a lateral intrusion of a vehicle's traffic.

As illustrated in FIG. 26, the memory 140 may store a lane restriction module 2610, an intruding vehicle acquisition module 2620, and an action module 2630.

[0342]

In one embodiment, the lane restriction module 2610 may store instructions, and when executed by the processing unit 110, the instructions may detect and define the first lane restriction 2502 and the second lane restriction 2504.

For example, the processing unit 110 may execute the lane offset restriction module 2610 to process a plurality of images received from at least one image capturing device 122-126 and detect the first lane restriction 2502 and the second lane restriction 2504. As discussed above, this may include identifying painted lane lines and/or measuring the midpoint of the road surface. The processing unit 110 may perform this analysis based on the technique described above in conjunction with FIGS. 5A to 5D.

[0343]

In one embodiment, the intruding vehicle acquisition module 2620 may store an instruction, and when executed by the processing unit 110, the instruction may detect the existence of the intruding vehicle 2510 and monitor the intruding vehicle 2510.

For example, the intrusion vehicle acquisition module 2620 may include instructions for determining the relative location of the intrusion vehicle 2510 with respect to the vehicle 200. In some embodiments, this may include monitoring the speed of the intruding vehicle 200. According to some embodiments, the intruding vehicle acquisition module 2620 may determine whether the intruding vehicle 2510 has exceeded the collision threshold. For example, the collision threshold may be the minimum predetermined distance between the vehicle 200 and the invading vehicle 2510. If the intruding vehicle 2510 is closer to the vehicle 200 than the minimum predetermined distance, the processing unit 110 may determine that the intruding vehicle 2510 is indeed intruding and execute an instruction to avoid a collision with the intruding vehicle 2510.

[0344]

Additionally or alternatively, the collision threshold may include the collision time.

The time-based collision threshold may be determined based on the following: the distance between the vehicle 200 and the intruding vehicle 2510, the speed of the vehicle 200 and/or the intruding vehicle 2510, the lateral speed of the vehicle 200 and/or the intruding vehicle 2510, and/or the vehicle 200 And/or the acceleration of the intruding vehicle 2510. The collision threshold based on the collision time may be set to allow sufficient time for the processing unit 110 and the system 100 to effectively perform an avoidance action to avoid a collision.

[0345]

Further, the processing unit 110 may execute the action module 2630 to respond to the detection of the intruding vehicle 2510 and/or to cross the collision threshold.

For example, the action module 2630 may include instructions for causing the vehicle 200 to maintain the current speed and travel within the first lane constraint 2502 and the second lane constraint 2504 such that the first distance 2530 is greater than the second distance 2540. The first distance 2530 may be located on the same side of the vehicle 200 as the intruding vehicle 2510, so that the vehicle 200 is further away from the intruding vehicle 2510 while still within the first lane constraint 2502 and the second lane constraint 2504. If the processing unit 110 determines that the collision threshold has been crossed, the

processing unit 110 may execute the action module 2630 to cause the vehicle 200 to take an evasive action, such as changing a route, braking, and/or changing lanes. The action module 2630 may also include instructions that cause the vehicle 200 to resume its course after completing the evasive action. For example, after braking or returning to the center of the lane 2505, the processing unit 110 may cause the vehicle 200 to restore the preset speed.

[0346]

FIG. 27 illustrates a flowchart of an exemplary process 2720 of detecting and responding to lateral intrusion of traffic by a vehicle consistent with the disclosed embodiments.

The process 2720 may identify the lane constraints that define the lane in which the vehicle travels, determine whether the intruding vehicle 2510 is approaching, and cause the vehicle 200 to maintain the current speed and travel within the first lane constraint 2501 and the second lane constraint 2504, so that the intrusion of the vehicle 200 The first distance 2530 on the side where the vehicle 2510 approaches is greater than the second distance 2540.

[0347]

In step 2730, at least one image capturing device 122, 124, and/or 126 may acquire multiple images of the area near the vehicle 200.

For example, a camera (such as the image capturing device 122) included in the image acquisition unit 120 may capture a plurality of images and send the images to the processing unit 110 via a digital connection (eg, USB, wireless, Bluetooth, etc.). In some embodiments, the processing unit 110 may receive more than one multiple images via multiple data interfaces. For example, the processing unit 110 may receive multiple images from each of the image capture devices 122, 124, 126, and each of the image capture devices 122, 124, 126 may have an associated data interface for communicating with the processing unit 110 data. The disclosed embodiments are not limited to any specific data interface structure or protocol.

[0348]

Next, in step 2740, the processing unit 110 may use the multiple images received via the data interface 128 to determine the first lane constraint on the first vehicle side 802 and the second lane constraint on the second vehicle side 812.

For example, as part of implementing step 2740, the processing unit 110 may execute instructions of the lane restriction module 2610. Further, as a part of determining the first lane constraint and the second lane constraint, the processing unit 110 may use one or more of the processes discussed above in conjunction with FIGS. 5A to 5D.

[0349]

In step 2750, the processing unit 110 may determine whether the intruding vehicle 2510 is approaching.

To make this determination, the processing unit 110 may execute the intrusion vehicle acquisition module 2620. The intrusive vehicle acquisition module 2620 may include instructions for processing images to detect other traffic. This may include identifying vehicles within a specific range of vehicle 200, such as those vehicles that are adjacent to vehicle 200. These neighboring vehicles may be invading vehicles 2510.

[0350]

Once the intruding vehicle 2510 is identified, the processing unit can determine more information about the vehicle 2510.

For example, in step 2555, the processing unit 110 may determine whether the target vehicle 2330 has exceeded the collision threshold. For example, the collision threshold may be the minimum predetermined distance between the vehicle 200 and the invading vehicle 2510. If the intruding vehicle 2510 is closer to the vehicle 200 than the minimum predetermined distance, the processing unit 110 may determine that the intruding vehicle 2510 is indeed intruding and execute an instruction to avoid a collision with the intruding vehicle 2510.

[0351]

Additionally or alternatively, the collision threshold may be the collision time.

IPR2025-00943
Tesla EX1002 Page 1234

The time-based collision threshold may be determined based on the following: the distance between the vehicle 200 and the intruding vehicle 2510, the speed of the vehicle 200 and/or the intruding vehicle 2510, the lateral speed of the vehicle 200 and/or the intruding vehicle 2510, and/or the vehicle 200 And/or the acceleration of the intruding vehicle 2510. The collision threshold based on the collision time may be set to allow sufficient time for the processing unit 110 and the system 100 to effectively perform an avoidance action to avoid a collision.

[0352]

In step 2760, the processing unit 110 may cause the vehicle 200 to maintain the current speed and travel within the first lane constraint 2502 and the second lane constraint 2504.

In order to cause the vehicle 200 to operate in this manner, the processing unit 110 may execute the action module 2630. Additionally or alternatively, step 2760 may include executing instructions on the action module 1230 to cause the processing unit 110 to send electronic signals to the accelerator 2610, brakes 2620, and/or steering system of the vehicle 200, so that the vehicle 200 maintains the current speed and is in the first position. The one-lane restriction 810 and the second-lane restriction 820 travel within.

[0353]

In step 2755, if the processing unit 110 determines that the collision threshold has been exceeded, the processing unit 110 may take an evasive action.

The evasive action may include causing the vehicle 200 to change lanes. A response may be expected, for example, that the vehicle 200 is traveling on a multi-lane highway, and the system 100 determines that changing lanes away from the invading vehicle 2510 will avoid a collision. The avoidance action may include causing the vehicle 200 to change its course. Additionally or alternatively, the avoidance action may include changing the route of the vehicle 200. For example, the processing unit 110 may take an evasive action by sending electronic signals to the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200. After completing the avoidance action, in step 2780, the processing unit 110 may control the vehicle 200 to resume its original navigation (for example, resume navigation to a desired location).

[0354]

Multi-threshold response zone for vehicle navigation

[0355]

The system 100 may provide driver assistance functionality for generating vehicle 200 responses, such as braking, accelerating, switching lanes, turning, and/or other navigation responses.

For example, if the system 100 detects that there is another vehicle in front of the vehicle 200 that is traveling at a slower speed than the vehicle 200, the system 100 may cause the vehicle 200 to reduce its speed by braking.

The system 100 may apply a braking profile that takes into account the immediacy of the risk (for example, the risk of collision) to provide the driver with a natural driving feeling.

[0356]

For example, the system 100 may generate a variety of different responses depending on the situation of the collision with another vehicle and the deterministic risk.

If the target vehicle is far ahead and the deterministic risk of collision is low, the system 100 may take no action in response to the speed or direction change of the target vehicle ahead. When the target vehicle may be closer or when the risk of collision rises above a predetermined level, the system 100 may generate a response, such as braking, and so on. When the risk of collision is even higher (for example, the system 100 determines that the collision is imminent without evasive action), the system 100 may cause the vehicle 100 to take an evasive action, such as maximum or near maximum braking, change of direction, and so on.

[0357]

28A and 28B are schematic representations of the host vehicle 200a on the road 2800 consistent with the disclosed embodiment.

In some examples, road 2800 may be divided into two lanes, lane 2810 and lane 2820. The target object, such as the target vehicle 200b, may travel in front of the host vehicle 200a in the same lane (for example, the lane 2820 shown in FIGS. 28A and 28B) or in a different lane. As shown in the example depicted in FIG. 28A, the host vehicle 200a and the target vehicle 200b may be separated by a distance d1. In a given time after the scene is depicted in FIG. 28A, as shown in FIG. 28B, the host vehicle 200a and the target vehicle 200b can be separated by a distance d2, where d2 is less than d1. For example, the host vehicle 200a can travel at a speed greater than that of the target vehicle 200b, causing the distance between the vehicles to decrease from d1 to d2.

[0358]

As another example, the host vehicle 200a and the target vehicle 200b can travel at the same speed until the target vehicle 200b reduces its speed through braking, or the host vehicle 200a increases its speed through acceleration, causing the distance between the vehicles to decrease from d1 to d2 .

As another example, the target vehicle 200b may intrude from one side of the host vehicle 200a. In some situations, if the target vehicle 200b is in front of the host vehicle 200a or far enough away from the host vehicle 200a (eg, five meters or more, 10 meters or more, etc.) and moves toward the host vehicle 200a, the system 100 may not take any action. However, if the target vehicle 200b continues to move closer to the host vehicle 200a, the system 100 may cause the host vehicle 200a to brake, change speed, and/or change lanes. For example, if the target vehicle 200b is within a predetermined threshold (e.g., within five meters), the host vehicle 200a may take action. There may be different scenarios involving the relative speed and position of the host vehicle 200a and the target vehicle 200b, and the disclosed embodiments are not limited to any specific scenario.

[0359]

In some embodiments, the multi-threshold response zone can provide the user with a natural-feeling driving experience.

For example, as discussed above, if the risk of collision is low, the system 100 may ignore the movement of the target vehicle 200b. As a result, when there is no risk of collision (for example, when the target vehicle 200b is 100 meters before, the host vehicle 200a will not need to imitate every movement of the target vehicle 200b), the host vehicle 200a may not respond to the movement of the target vehicle 200b or action. However, if the target vehicle 200b is close to the host vehicle 200a and/or if the system 100 determines that the risk of collision is high, a responsive maneuver can be taken to make the user feel safer by increasing the space between the host vehicle 200a and the target vehicle 200b. If the system 100 determines that an impact may be imminent, the system 100 can take a more aggressive action without surprising the user, as this action may be necessary to avoid the impact.

[0360]

Therefore, in some embodiments, the system 100 may associate different reaction zones with different degrees of response actions.

For example, in some embodiments, when the target vehicle 200b is within the first threshold (e.g., when the target vehicle 200b is far away from the host vehicle 200a), the system 100 may take no action; when the target vehicle 200b is within the second threshold (e.g., , When the target vehicle 200b is closer to the host vehicle 200a), the system 100 can take a moderate action; and when the target vehicle 200 is within the third threshold (for example, when the target vehicle 200b is close enough to the host vehicle 200a, there is a risk of collision), the system 100 can take more radical actions. In the first reaction zone associated with the first threshold, when the estimated distance from the target associated with the target vehicle 200b is more than twice the predetermined safe distance, the system 100 may take no action. In some embodiments, the first reaction zone may be applied in the case where the estimated distance from the target associated with the target vehicle 200b falls within the range of 1.5 times to 2.5 times the predetermined safety distance.

[0361]

The estimated distance to the target may be the current distance between the host vehicle 200a and the target vehicle 200b plus the predicted time of the distance (for example, in the range of 0.1 second to 1.0 second, depending on factors associated with the system 100, such as speed Track the expected changes after the control loop, actuators, and vehicle dynamics.

Tesla EX1002 Page 1236 IPR2025-00943

The expected distance change can be calculated as: (v target-v main) * t prediction + ((a target-a main) * t prediction 1/2), where v target and v main correspond to target vehicle 200b and host vehicle 200a respectively The speeds of a target and a main correspond to the accelerations of the target vehicle 200b and the host vehicle 200a, respectively, and the t prediction corresponds to the predicted time. The selection of the predicted time that approximates the overall reaction time of the system 100 may take into consideration the fact that the desired navigation response of the host vehicle 200a may not be implemented immediately, so as to provide the driver and/or passengers of the host vehicle 200a with a smooth and natural driving experience.

[0362]

The predetermined safety distance can be defined as: maximum {d static, d dynamic}.

Dstatic, which may fall within a range of 2 to 10 meters, may indicate a desired safety distance between the host vehicle 200a and another vehicle (such as the target vehicle 200b) when the vehicle is stopped and/or is moving at a low speed. d dynamic can represent the expected safety distance when the vehicle is moving at a speed greater than low speed, and can be calculated as: t safety * minimum (v target, v main), where v target and v main correspond to target vehicle 200b and host vehicle 200a, respectively speed. t safety may fall within the range of 0.5 seconds to 2.5 seconds, and may be adjusted by the driver of the host vehicle 200a via the user interface 170 according to the driver's preference. Thus, the driver can control the distance maintained by the system 100 between the host vehicle 200a and other vehicles.

[0363]

In the second reaction zone associated with the second threshold, the system 100 may place the target vehicle 200b in the safe zone and apply different weights to different maneuvers performed by the target vehicle 200b.

For example, when the target vehicle 200b can approach the lower boundary of the safe zone, for example, the estimated distance from the target associated with the target vehicle 200b can be close to the predetermined safe distance (discussed above in conjunction with the first reaction zone), and the system 100 can decelerate to maneuver. Assign more weight (e.g., imitate more) and assign a smaller weight to maneuver acceleration (e.g., imitate less). In some embodiments, the lower boundary may fall within the range of 0.5 times to 1.5 times the predetermined safety distance (depending on the driver's preference). Conversely, when the target vehicle 200b approaches the upper boundary of the safety zone, for example, the estimated distance from the target associated with the target vehicle 200b may be close to twice the predetermined safety distance, and the system 100 may assign a smaller weight to the maneuvering deceleration (for example, , imitate less) and assign more weight to maneuver acceleration (for example, imitate more). In some embodiments, the lower boundary may fall within the range of 1.5 times to 2.5 times the predetermined safety distance (depending on the driver's preference).

[0364]

In the third reaction zone associated with the third threshold, when the estimated distance from the target associated with the target vehicle 200b is less than a predetermined safe distance, the system 100 may place the target vehicle 200b in the danger zone.

In this scenario, the system 100 may cause the host vehicle 200a to imitate (eg, immediately imitate) any deceleration maneuvers performed by the target vehicle 200b (while ignoring any acceleration maneuvers). In some embodiments, in the case where the estimated distance from the target associated with the target vehicle 200b falls within the range of 0.5 to 1.5 times the predetermined safety distance (depending on the driver's preference), the first reaction zone may be applied .

[0365]

In some embodiments, the system 100 may generate a response of the host vehicle 200a, such as braking.

As indicated in FIG. 28B, for example, the distance between the host vehicle 200a and the target vehicle 200b may be reduced from d1 to d2. Depending on the distance d2 and extending to the time before the collision, the system 100 may cause the host vehicle 200a to reduce its speed to avoid a collision with the target vehicle 200b. Additionally, based on, for example, the time before the collision, the system 100 may apply one or more braking profiles when the host vehicle 200a is caused to reduce its speed. Thus, if the time (and/or distance d2) before the collision exceeds the first predetermined threshold, the system 100 may cause the host vehicle 200a to brake, change lanes, and/or change speed in a gradual manner, and if the time before the

Tesla EX1002 Page 1237 IPR2025-00943

collision When the time (and/or the distance d2) is lower than the second predetermined threshold, the system 100 may cause the host vehicle 200a to brake, change lanes, and/or change speed in a faster manner.

[0366]

Figure 29 is an exemplary block diagram of a memory consistent with the disclosed embodiments, the memory being configured to store instructions for performing one or more operations.

As shown in FIG. 29, the memory 140 can store a target monitoring module 2902, an interception time module 2904, and an action response module 2906. The disclosed embodiments are not limited to the memory 140 of any particular configuration. Further, the processing unit 110 can execute instructions stored in any of the modules 2902-2906 (included in the memory 140).

[0367]

In one embodiment, the target monitoring module 2902 may store software instructions (such as computer vision software) that, when executed by the processing unit 110, identify and monitor targets appearing in multiple images representing the environment near the host vehicle 200 Object (such as target vehicle 200b).

For the purpose of the present disclosure, the target object may also be referred to as the target vehicle 200b, but the target object may be pedestrians, road hazards (eg, large debris), and so on. The multiple images may be acquired by one or more of the image capture devices 122, 124, and 126. The processing unit 110 may execute the target monitoring module 2902 to identify and monitor the target vehicle 200b by, for example, analyzing multiple images using the technology described above in conjunction with FIGS. 5A to 5D. Thus, the target monitoring module 2902 can track different variables, such as the motion (e.g., direction and trajectory), speed and acceleration associated with the target vehicle 200b, and the distance between the host vehicle 200a and the target vehicle 200b. Such information can be tracked continuously or at predetermined intervals.

[0368]

In one embodiment, the interception time module 2904 may store software instructions that, when executed by the processing unit 110, determine an indicator of the interception time between the host vehicle 200a and the target vehicle 200b.

The interception time may refer to the amount of time before the host vehicle 200a comes into contact (eg, collides) with the target vehicle 200b. The processing unit 110 may execute the interception time module 2904 to determine an indicator of the interception time based on analyzing images acquired by one or more of the image capturing devices 122, 124, and 126. For example, the intercept time module 2904 may use a series of time-based observations (such as Kalman filter or linear quadratic estimation (LQE)) of the position, velocity, and/or acceleration of the target vehicle 200b (relative to the host vehicle 200a), for example. Make sure. The Kalman filter may be based on a measurement result of the scale of the target vehicle 200b, where the measurement result of the scale is proportional to the interception time.

[0369]

In one embodiment, the action response module 2906 may store software instructions (such as computer vision software) that, when executed by the processing unit 110, generate one or more responses in the host vehicle 200a.

The response may be based on the identification and monitoring of the target vehicle 200b (e.g., performed via the execution target monitoring module 2902) and/or an indicator of the interception time between the host vehicle 200a and the target vehicle 200b (e.g., via the execution interception time module 2904 confirmed). The processing unit 110 may generate one or more responses in the host vehicle 200a by sending electronic signals to the throttle system 220, the braking system 230, and/or the steering system 240 of the host vehicle 200a, for example (for example, via the CAN bus) to Braking, accelerating, or triggering a turn, lane shift, and/or direction change of the host vehicle 200a. The processing unit 110 may use the extended Kalman filter to estimate data associated with the target vehicle 200b and the host vehicle 200a, such as speed, acceleration, turning rate, pitch angle, pitch rate, and the position of the target vehicle 200b relative to the host vehicle 200a And heading, size of target vehicle 200b, road geometry, etc. For example, the processing unit 110 may receive measurement results from components on the host vehicle 200a (such as an odometer, an inertial measurement unit (IMU), an image processor 190, etc.). Thus, the processing unit 110 can use the state formulas, rigid constraints, and motion assumptions derived from the kinematic formulas of motion to calculate data associated with the target vehicle 200b and/or the host vehicle 200a, where the input of the formula may include the input from the host vehicle 200a The measurement results of the components on the. The processing unit 110

IPR2025-00943 Tesla EX1002 Page 1238

may also use the techniques described above in conjunction with FIGS. 4 to 7 to generate one or more responses of the host vehicle 200a.

[0370]

FIG. 30 is a flowchart illustrating an exemplary process 3000 for generating a response of the host vehicle consistent with the disclosed embodiment.

In step 3010, the processing unit 110 may receive a plurality of images via the data interface 128 located between the processing unit 110 and the image acquisition unit 120. For example, one or more cameras included in the image acquisition unit 120 can capture multiple images of the area in front of the host vehicle 200a (for example, or the side or back of the vehicle) and connect via data (for example, digital, wired, USB , Wireless, Bluetooth, etc.) send the image to the processing unit 110.

[0371]

In step 3020, the processing unit 110 may execute the target monitoring module 2902 to identify target objects (such as the target vehicle 200b) in a plurality of images.

For example, the target monitoring module 2902 may use the technique described above in conjunction with FIG. 5A and FIG. 5B to detect the presence of a vehicle in the image. Based on image data, location data (for example, GPS location information), map data, the yaw rate of the host vehicle 200a, lane or other road markings, speed data, and/or data from sensors included in the host vehicle 200a, multiple vehicles (And/or the object) may be detected in the same lane as the host vehicle 200a or a different lane, and the distance from the host vehicle 200a is variable. In the scenario of detecting multiple vehicles, the target monitoring module 2902 may determine the target vehicle 200b as the vehicle traveling in the same lane as the host vehicle 200a and/or the vehicle closest to the host vehicle 200a.

[0372]

In step 3030, the processing unit 110 can execute the target monitoring module 2902 to monitor the target vehicle 200b identified in step 3020.

For example, the target monitoring module 2902 may track information associated with the target vehicle 200b, such as movement (e.g., direction of travel and trajectory), speed and acceleration, and the distance between the host vehicle 200a and the target vehicle 200b. The target monitoring module 2902 can use the technique described above in conjunction with FIGS. 5A to 5D to track such information continuously or at predetermined intervals by analyzing multiple images. In step 3040, the processing unit 110 may execute the interception time module 2904 to determine an indicator of the interception time between the host vehicle 200a and the target vehicle 200b. The interception time module 2904 may make the determination by tracking the position, speed, and/or acceleration of the target vehicle 200b (relative to the host vehicle 200a) over one or more time periods. The interception time module 2904 may use the techniques described above in conjunction with FIGS. 5A and 5B to track such information.

[0373]

In step 3050, the processing unit 110 may execute the action response module 2906 to generate one or more responses of the host vehicle 200a.

For example, the processing unit 110 may (for example, via the CAN bus) send electronic signals to the throttle system 220, the braking system 230, and/or the steering system 240 of the host vehicle 200a to brake, accelerate or trigger the host vehicle 200a to turn, Lane shift and/or direction change. Further, one or more actuators may control the throttle system 220, the braking system 230, and/or the steering system 240. For example, the processing unit 110 may send an electronic signal to cause the system 100 to physically depress the brake for a predetermined amount or partially release the accelerator of the vehicle 200. Further, the processing unit 110 may send an electronic signal that causes the vehicle 200 of the system 100 to turn in a specific direction. Such a response may be based on the monitoring of the target vehicle 200b (performed at step 3030) and/or an indicator of the interception time between the host vehicle 200a and the target vehicle 200b (determined at step 3040). In some embodiments, the interception time may be compared with one or more predetermined thresholds (e.g., first and second interception thresholds). For example, if the indicator of interception time indicates that the target object is within a certain interception threshold (for example, within a second interception threshold), the response may include an emergency avoidance action. Emergency avoidance actions may include lane changes and/or emergency braking.

IPR2025-00943
Tesla EX1002 Page 1239

[0374]

The threshold can be expressed as a unit of time.

Thus, for example, if the indicator of the interception time indicates that the target vehicle 200b is outside the first threshold (for example, the time before the collision between the host vehicle 200a and the target vehicle 200b exceeds the first threshold), the action response module 2906 The response may be generated in the host vehicle 200a based on the average motion profile associated with the target vehicle 200b. The average motion profile associated with the target vehicle 200b may represent the average position and/or average speed of the target vehicle 200b with respect to a certain threshold. In some embodiments, if the target vehicle 200b is a certain distance and/or is traveling at a certain speed, the system 100 may take no action. For example, if the target vehicle 200b is in front of the host vehicle 200a or is far enough away from the host vehicle 200a (eg, five meters, 10 meters, etc.), the system 100 may take no action unless the target vehicle 200b is closer to the host vehicle 200.

[0375]

If the indicator of the interception time indicates that the target vehicle 200b is between the first threshold and the second threshold (for example, the time before the collision between the host vehicle 200a and the target vehicle 200b falls between the first threshold and the second threshold) , The action response module 2906 may generate the response of the host vehicle 200a based on the actual motion associated with the target vehicle 200b.

The actual movement may represent the position and/or speed of the target vehicle 200b (for example, not averaged or delayed), if the interception time indicator indicates that the target vehicle 200b is within the second threshold (for example, between the host vehicle 200a and the target vehicle 200b) The time before the collision is lower than the second threshold), the action response module 2906 may generate an emergency action (such as lane change or emergency braking) of the host vehicle 200a.

[0376]

FIG. 31 is a flowchart showing an exemplary process 3100 for refusing to generate a response in the host vehicle consistent with the disclosed embodiment.

In step 3110, as described above in conjunction with step 3010 of FIG. 30, the processing unit 110 may receive multiple images via the data interface 128 located between the processing unit 110 and the image acquisition unit 120.

[0377]

In step 3120, as described above in connection with step 3020 of FIG. 30, the processing unit 110 may execute the target monitoring module 2902 to locate the target object (for example, the target vehicle 200b) in the image received in step 3010 by analyzing the image.

In step 3125, the processing unit 110 may execute the target monitoring module 2902 to determine whether the target object (for example, the target vehicle 200b) is traveling in a lane different from the host vehicle 200a using the technique described above in conjunction with the step 3030 of FIG. 30. If it is determined that the target vehicle 200b is traveling in a lane different from the host vehicle 200a (step 3125, YES), then in step 3130, the processing unit 110 may refuse to generate a response in the host vehicle 200a. Alternatively, if the system 100 determines that the target vehicle 200b does not pose an imminent risk of collision (for example, the target vehicle 200b is on a non-intersecting trajectory with the host vehicle 200a or the potential time to collide exceeds the threshold, so that the target object is at the current time It does not constitute an imminent risk), in step 3130, the processing unit 110 may refuse to generate a response in the host vehicle 200a. Otherwise (step 3125, No), the process 3100 can end without a response.

[0378]

Use a multi-section braking profile to stop autonomously driving vehicles

[0379]

The system 100 may provide driver assistance functionality for monitoring and responding to environmental conditions to control the braking of the vehicle 200.

For example, the system 100 may detect an object that triggers the stop of the vehicle 200 based on multiple images.

For example, the object may include another vehicle, a traffic light, and/or a traffic sign near the vehicle 200. The system 100 may use the braking profile to slow down and/or stop the vehicle 200. According to some embodiments, the braking profile may include multiple sections that perform deceleration at different stages. For example, the braking profile may cause the system 100 to stop the vehicle 200 at a reduced deceleration.

[0380]

By providing a braking profile that includes multiple segments, the system 100 can cause the vehicle 200 to operate in a manner that feels natural to the driver and/or passengers.

For example, when the driver stops the vehicle at a traffic light, the brake is usually not applied at a constant speed from the moment of application until the moment the vehicle stops. Operating the brake in this way results in a sudden stop at termination. Conversely, the driver can first apply the brake gently, then increase the pressure, and then when the vehicle reaches its target stop, the driver can reduce the braking level until it reaches zero (or near zero) when the vehicle stops. Operating the vehicle in this way can provide a natural stopping process, that is, the driver can first cause a gradual deceleration, so the vehicle will not experience a sudden initial deceleration and then terminate, so that the vehicle will not stop suddenly.

[0381]

The system 100 may provide autonomous vehicle navigation that may implement braking according to a method that feels natural to the driver and/or passengers of the vehicle 200.

For example, in some embodiments, the system 100 may cause one or more of the image capture devices 122-126 to acquire multiple images of the area near the vehicle 200. The processing device 110 may receive a plurality of images via the data interface 128, and identify triggers (for example, objects, traffic lights, and/or traffic signs) for stopping the vehicle 200 based on the analysis of the plurality of images. Based on the identified trigger, the system 100 can cause the vehicle 200 to stop according to the braking profile. In some embodiments, as determined based on an analysis of a plurality of images, the braking profile may include a plurality of segments (for example, including: a first segment associated with a first deceleration rate; including a first segment associated with a first deceleration rate; The second section of the second deceleration rate; and the third section where the braking level is reduced when approaching the target stopping point).

[0382]

FIG. 32A illustrates a vehicle 200 traveling on a road 3210 including an object 3220.

The system 100 may identify a trigger (e.g., object 3220) to stop the vehicle 200 relative to the detection of the trigger. For example, the system 100 may first decelerate the vehicle 200 after the trigger is detected, and stop the vehicle 200 before reaching the object 3220.

[0383]

The processing unit 110 may be configured to determine a trigger for stopping the vehicle 200 based on a plurality of images acquired by one or more of the image capturing devices 122-126.

The processing unit 110 may receive images from the image capture devices 122-126 via the data interface 128. According to some embodiments, the trigger may include an object 3220. For example, the object 3220 may be another vehicle on the road 2810. According to some embodiments, if the object 3220 is in a certain nearby area of the vehicle 200, the processing unit 110 may identify the trigger. For example, the recognition trigger may depend on the relative distance between the object 3220 and the vehicle 200. Further, this identification optionally considers the relative speed and/or relative acceleration between the object 3220 and the vehicle 200.

[0384]

Tesla EX1002 Page 1241 IPR2025-00943

Additionally or alternatively, the trigger may include a traffic sign, such as the traffic sign 3240 shown in FIG. 32A.

For example, the traffic sign 3240 may be a stop sign and/or a speed limit sign. According to some embodiments, the processing unit 110 may process multiple images to determine the content of the traffic sign 3240 as part of determining whether the trigger exists. For example, if the vehicle 200 exceeds the speed limit, the processing unit 110 may determine the existence of the trigger based on the speed limit specified by the traffic sign 3240. As another example, if the yellow light is on and/or the red light is on, the processing unit 110 may determine the existence of the trigger based on the traffic light 2830.

[0385]

Based on the multiple images, the processing unit 110 may determine the target stopping location 3250, and once the braking profile is executed, the vehicle 200 will stop at the target stopping location 3250.

Reaching the target stopping point 3250 may be accompanied by the release of the brake pressure of the vehicle 200. The target stop location 3250 may be based on the location of the traffic light 3230 and/or the traffic sign 3240 such that the vehicle 200 conforms to the traffic light 3230 and/or the traffic sign 3240. Alternatively, the target stop location 3250 may be based on the object 3220 such that the vehicle 200 stops before reaching the object 3220. Therefore, the target stop location 3250 may be based on the location of the object 3220 to avoid a collision between the vehicle 200 and the object 3220. In some embodiments, the update of the target stop location 3250 may be based on new information obtained from multiple images, such as newly identified targets, new visible obstacles, new visible signs, and so on.

[0386]

Figure 32B illustrates an exemplary braking profile consistent with the disclosed embodiment.

In the example shown in Figure 32B, the braking profile includes four sections. However, the number of sections is exemplary, and a braking profile consistent with the disclosed embodiments may include any suitable number of sections (eg, 2, 3, 4, 5, 6, etc. Section).

[0387]

In section 1, the system 100 may prepare for possible braking scenarios.

For example, in this area, the system 100 may receive information (e.g., map data and/or image data) about an upcoming intersection (e.g., an intersection with traffic lights). The system 100 may gradually reduce the speed of the vehicle 200 (for example, depending on the characteristics of the intersection, the speed of the vehicle 200 may be gradually reduced to about 70-50 km/hr). The system 100 can generate speed adjustments, for example, about an average deceleration of about 0.2-0.4 m/s2 (approximately 0.5 m/s2 at the maximum) about 100-200 meters before the hub. By reducing the speed of the vehicle 200 in section 1, the system 100 can provide the driver and/or passengers of the vehicle 200 with the confidence that the system 100 knows the approach to the intersection.

[0388]

In section 2, the system 100 can apply strong braking.

In this section, the system 100 may receive information based on the analysis of the image data, for example, a traffic light turns red and/or there is a stopped vehicle at an intersection in the travel line of the vehicle 200. The system 100 can cause the vehicle 200 to experience a significant speed reduction in this zone, while the distance from the stop line or stopped vehicle is still large. Doing so can provide a comfortable feeling for the driver and/or passengers of the vehicle 200, and the system 100 will have enough space to complete the braking maneuver. The zone may be about 30-100 meters away from the stop line, and the average deceleration may be about 1.5-2.5 m/s2 (the maximum is about 3.5 m/s2).

[0389]

In section 3, the system 100 can make moderate braking adjustments.

For example, in this section, the system 100 may adjust the speed of the vehicle 200 according to the remaining distance from the stop line or the stopped vehicle and based on the current speed and deceleration of the vehicle 200. By making this adjustment, the system 100 can provide the driver and/or passenger of the vehicle 200 with an indication that the system 100 is

releasing most of the braking power while the speed is still decreasing, but slower than before. The zone can be about 5-30 meters away from the stop line or stop the vehicle, and the average deceleration can be about 0.5-1.5 m/s2 (the maximum is about 2 m/s2).

[0390]

In section 4, the system 100 can make small braking adjustments.

In this section, the system 100 can shorten the remaining distance from the stop line or the stopped vehicle at a very low speed gradually becoming zero. In this zone, the driver and/or passenger may feel the vehicle 200 slowly sliding into place (e.g., at a stop line or behind another vehicle). This zone may constitute about the last 5-7 meters from the stop line or stopped vehicle, and the average deceleration may be about 0.3-0.5 m/s2 (the maximum is about 1 m/s2).

[0391]

FIG. 33 is an exemplary block diagram of the memory 140 and/or 150 consistent with the disclosed embodiments. The memory 140 and/or 150 may store instructions for performing one or more operations.

As illustrated in FIG. 33, the memory 140 may store one or more modules for performing trigger detection and response described herein. For example, the memory 140 may store the trigger recognition module 3300 and the action response module 3310.

[0392]

The trigger recognition module 3300 may store an instruction, and when executed by the processing unit 110, the instruction may detect the existence of a trigger for stopping the vehicle 200.

For example, the trigger recognition module 3300 may process a plurality of images received from at least one image capture device 122-124 to detect the presence of a trigger. As discussed above, this may include identifying objects 3220, traffic lights 3230, and/or traffic signs 3240.

[0393]

The action response module 3310 may store an instruction, and when executed by the processing unit 110, the instruction may respond to the existence of a trigger.

For example, the action response module 3310 may perform control to decelerate and/or stop the vehicle 200. For example, the action response module 3310 may perform a braking profile to slow down and/or stop the vehicle 200. The action response module 3310 may determine to execute one of a plurality of braking profiles based on the number of braking profiles and based on, for example, environmental factors, the type of trigger being detected, and the speed of the vehicle 200. In order to navigate the vehicle 200 to a stop based on the detected trigger, the processing unit 110 may send an electronic signal to one or more of the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200. Further, one or more actuators may control the throttle system 220, the braking system 230, and/or the steering system 240. For example, the processing unit 110 may send an electronic signal to cause the system 100 to physically depress the brake for a predetermined amount.

[0394]

FIG. 34 illustrates a process 3400 for navigating the vehicle 200 consistent with the disclosed embodiment.

The process 3000 may identify a trigger for stopping the vehicle 200 (eg, an object, vehicle, traffic light, traffic sign, etc.) and execute a braking profile in response to the detection of the trigger. The process 3400 may detect various types of triggers to determine which of the multiple braking profiles to perform.

[0395]

The braking profile may be based on the type of the detected trigger and/or other environmental conditions related to the detected trigger (eg, distance from the object).

IPR2025-00943
Tesla EX1002 Page 1243

For example, in some embodiments, the braking profile may include multiple sections (e.g., three sections). In the three-section braking profile, the first section of the braking profile may be associated with a first deceleration rate, the second section of the braking profile may include a second deceleration rate that is less than the first deceleration rate, and the braking profile The third section of may include causing the braking level to decrease when approaching the target stopping location, which is determined based on an analysis of multiple images acquired by one or more of the image capture devices 122-126. For example, in some embodiments, the first section may cause the vehicle 200 to gradually increase deceleration, the second section may cause the vehicle 200 to decelerate at a constant speed, and the third section may cause the vehicle 200 to gradually reduce its deceleration rate to zero. (Or near zero).

[0396]

In step 3410, the process 3400 may use at least one of the image capture devices 122-126 to acquire multiple images of the area near the vehicle 200.

For example, the processing unit 110 may receive multiple images via the data interface 128. The processing unit 110 may be configured to determine the presence of one or more objects and/or signs located on or near the road 3210 based on the acquired image. For example, the processing unit 110 may be configured to determine the existence of the object 3220, the traffic light 3230, and/or the traffic sign 3240. The processing unit 110 may be further configured to read the content of the traffic sign 3240. For example, the processing unit 110 may be configured to distinguish between a speed limit sign and a stop sign. Further, the processing unit 110 may be configured to process the image based on the traffic sign 3240 to determine the speed limit.

[0397]

At step 3420, the process 300 may include a trigger for stopping the vehicle 200.

For example, the processing unit 110 may identify the trigger based on the existence of the object 3220 in the area near the vehicle 200. Additionally or alternatively, the processing unit 110 may identify the trigger based on the current signal state of the traffic light 3230. According to some embodiments, if it is not a green light, the processing unit 110 may determine that the traffic light 3230 constitutes a trigger. According to some embodiments, for example, if the traffic sign 3240 is a stop sign, the processing unit 110 may determine that the traffic sign 3240 constitutes a trigger. In some embodiments, if the traffic sign 3240 is a speed limit sign and the vehicle 200 exceeds the speed limit, the traffic sign 3240 may constitute a trigger. Further in some embodiments, the processing unit 110 may cause the system 100 to provide a voice notification, recognize a trigger, and/or indicate that the vehicle 200 is braking or is about to start braking.

[0398]

After the processing unit 110 recognizes the trigger, in step 3430, the processing unit 110 may cause the vehicle 200 to decelerate and/or stop according to the braking profile.

In some embodiments, the braking profile may include a certain number of profile sections. The different contour sections of the braking contour can be defined by the time period and/or the vehicle speed. For example, the processing unit 110 may perform the control of the vehicle 200 according to the first contour section lasting a predetermined period of time or until the speed of the vehicle 200 reaches the target value. The target value may be a predetermined speed, or may be a percentage of the speed at which the vehicle 200 has traveled in the past before the braking profile was performed. In order to implement the braking profile, the processing unit 110 may send an electronic signal to one or more of the throttle system 220, the braking system 230, and/or the steering system 240 of the vehicle 200. Further, one or more actuators may control the throttle system 220, the braking system 230, and/or the steering system 240.

[0399]

As discussed above, according to some embodiments, the profile section of the braking profile may correspond to the deceleration rate.

For example, the braking profile may include multiple sections, and the first section may be associated with a first deceleration rate. The braking profile may further include a second section associated with a second deceleration rate. The second deceleration rate may be less than the first deceleration rate.

IPR2025-00943
Tesla EX1002 Page 1244

[0400]

Further, as discussed above, according to some embodiments, the braking profile may include a third profile section, where a reduction in the braking level is a potential factor for the vehicle 200 to generate triggers (e.g., object 3220, traffic light 3230, traffic Flag 3240) as a function of the distance between.

For example, the braking profile may include reducing the braking level as the vehicle 200 approaches the target stopping location 3250. Further, reaching the target stopping point 3250 may be accompanied by the release of the brake pressure of the vehicle 200.

[0401]

For illustration purposes, the foregoing description has been presented.

It is not exhaustive and not limited to the specific forms or implementations disclosed. Considering the practice of the specification and disclosed embodiments, modifications and rewriting will be obvious to those skilled in the art. In addition, although various aspects of the disclosed embodiments are described as being stored in a memory, those skilled in the art will understand that these aspects may also be stored on other types of computer-readable media such as secondary storage devices, such as hard disks. Or CD ROM, or other forms of RAM or ROM, USB media, DVD, Blu-ray or other optical drive media.

[0402]

Computer programs based on written descriptions and published methods are within the skills of experienced developers.

Various programs or program modules can be created using any technology known to those skilled in the art or can be designed in combination with existing software. For example, program segments or program modules can be designed or used. Net Framework,. Net Compact Framework (and related languages such as Visual Basic, C, etc.), Java, C++, Objective-C, HTML, HTML/AJAX combination, XML or HTML including Java applets.

[0403]

Moreover, although illustrative embodiments have been described herein, the scope of any and all embodiments having equivalent elements, modifications, omissions, combinations (for example, aspects across various embodiments), adaptations, and/or changes will be based on the present disclosure Known by those skilled in the art.

The limitations in the claims should be interpreted broadly based on the language used in the claims, and are not limited to the examples described in this specification or during the prosecution of this application. Examples should be interpreted as non-exclusive. In addition, the steps in the disclosed method can be modified in any way, including reordering the steps and/or inserting or deleting steps. Therefore, it is expected that the specification and examples are considered to be illustrative only, and the true scope and spirit are indicated by the following claims and their full equivalents.

[0404]

Cross-references to related applications

[0405]

This application claims the following priority: U.S. Provisional Patent Application No. 61/911,490 filed on December 4, 2013; U.S. Provisional Patent Application No. 61/993,084 filed on May 14, 2014; Filed on May 14, 2014 United States Provisional Patent Application No. 61/993,111; United States Provisional Patent Application 62/015,524 filed on June 23, 2014; United States Provisional Patent Application 62/022,221 filed on July 9, 2014; United States filed on August 21, 2014 Provisional Patent Application No. 62/040,224; and U.S. Provisional Patent Application 62/040,269 filed on August 21, 2014.

The entire contents of all aforementioned applications are incorporated herein by reference.

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

CLAIMS CN110920609A

1.

An autonomous vehicle control system, which includes:

A plurality of cameras, the plurality of cameras are used to obtain a plurality of images of the vicinity of the vehicle; and

At least one processing device configured to:

Identify the curve to be navigated based on map data and vehicle location information;

Determining the initial target speed of the vehicle based on at least one characteristic of the curve reflected in the map data;

Adjusting the speed of the vehicle to the initial target speed;

Determining the observed characteristic of the curve based on the plurality of images;

Determining an updated target speed based on the observed characteristics of the curve; and

The speed of the vehicle is adjusted to the updated target speed.

2.

An autonomous vehicle control system for controlling a vehicle, the autonomous vehicle control system comprising:

At least one processing device configured to:

Determining the initial target speed of the vehicle based on global positioning data;

Adjusting the speed of the vehicle to the initial target speed for the first part of the navigation event;

Determining an updated target speed based on sensor data, the updated target speed being used in the second part of the navigation event; and

The speed of the vehicle is adjusted to the updated target speed during the second part of the navigation event.

3.

The autonomous vehicle control system according to claim 2, wherein the navigation event is to navigate a curve in a road, and wherein the first part of the navigation event includes the vehicle before the curve Navigate on the road, and the second part of the navigation event includes navigating the vehicle on the road through the curve.

4.

The autonomous vehicle control system according to claim 2, wherein the initial target speed of the vehicle is determined based on global positioning data, and the at least one processing device recognizes the upcoming vehicle based on map data and the global positioning data Road section.

5.

The autonomous vehicle control system according to claim 2, wherein the sensor data includes image data.

6.

The autonomous vehicle control system according to claim 2, the autonomous vehicle control system includes a plurality of cameras to acquire a plurality of images of an area in the traveling direction of the vehicle.

7.

The autonomous vehicle control system according to claim 6, wherein the at least one processing device determines a road characteristic based on the plurality of images and uses the road characteristic to determine the updated target speed.

(19)中华人民共和国国家知识产权局

(12)发明专利申请

(10)申请公布号 CN 110920609 A
(43)申请公布日 2020.03.27

(21)申请号 201911164109.8

(22)申请日 2014.12.04

(30)优先权数据
61/911,490 2013.12.04 US
61/993,084 2014.05.14 US
61/993,111 2014.05.14 US
62/015,524 2014.06.23 US
62/022,221 2014.07.09 US
62/040,224 2014.08.21 US
62/040,269 2014.08.21 US

(62)分案原申请数据
201480074153.9 2014.12.04

(71)申请人 移动眼视力科技有限公司
地址 以色列耶路撒冷

(72)发明人 A·克拉克 E·宝葛 G·尼尔

(74)专利代理机构 北京三友知识产权代理有限公司 11127
代理人 王小东

(51)Int.Cl.
B60W 30/12(2020.01)
B60W 30/18(2012.01)
G01C 21/36(2006.01)
G01C 21/34(2006.01)
G06K 9/00(2006.01)

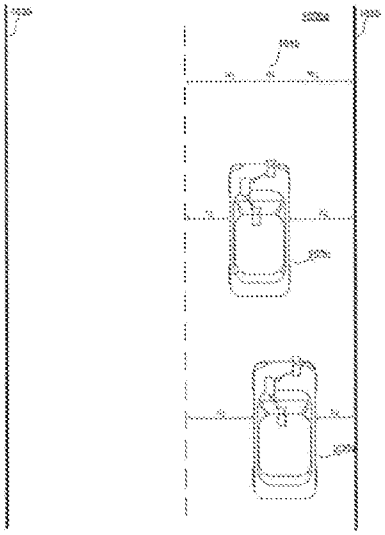
权利要求书1页 说明书62页 附图55页

(54)发明名称

用于模仿前车的系统和方法

(57)摘要

本发明涉及用于模仿前车的系统和方法。系统和方法使用相机来提供自主导航特征。在一个实施中，提供了一种用于主车辆的驾驶员辅助导航系统。所述系统可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取主车辆的附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：在所述多个图像中定位前车；基于所述多个图像来确定所述前车采取的至少一个动作；以及致使所述主车辆模仿所述前车的所述至少一个动作。



CN 110920609 A

IPR2025-00943
Tesla EX1002 Page 1248

1.一种自主车辆控制系统,所述自主车辆控制系统包括:

多个相机,所述多个相机用以获取车辆的附近区域的多个图像;以及

至少一个处理装置,所述至少一个处理装置被构造成:

基于地图数据和车辆位置信息来辨识待导航的弯道;

基于反映在所述地图数据中的所述弯道的至少一个特性来确定所述车辆的初始目标速度;

将所述车辆的速度调整至所述初始目标速度;

基于所述多个图像来确定所述弯道的观察到的特性;

基于所述弯道的所述观察到的特性来确定更新的目标速度;以及

将所述车辆的速度调整至所述更新的目标速度。

2.一种用于控制车辆的自主车辆控制系统,所述自主车辆控制系统包括:

至少一个处理装置,所述至少一个处理装置被构造成:

基于全球定位数据来确定所述车辆的初始目标速度;

针对导航事件的第一部分将所述车辆的速度向所述初始目标速度调整;

基于传感器数据来确定更新的目标速度,所述更新的目标速度用在所述导航事件的第二部分中;以及

在所述导航事件的所述第二部分期间将所述车辆的速度向所述更新的目标速度调整。

3.根据权利要求2所述的自主车辆控制系统,其中,所述导航事件是对于道路中的弯道进行导航,其中所述导航事件的所述第一部分包括使所述车辆在所述弯道之前的道路上导航,并且所述导航事件的所述第二部分包括使所述车辆在通过所述弯道的道路上导航。

4.根据权利要求2所述的自主车辆控制系统,其中,基于全球定位数据来确定所述车辆的所述初始目标速度,所述至少一个处理装置基于地图数据和所述全球定位数据来辨识即将到来的道路区段。

5.根据权利要求2所述的自主车辆控制系统,其中,所述传感器数据包括图像数据。

6.根据权利要求2所述的自主车辆控制系统,所述自主车辆控制系统包括用以获取所述车辆的行进方向上的区域的多个图像的多个相机。

7.根据权利要求6所述的自主车辆控制系统,其中,所述至少一个处理装置基于所述多个图像来确定道路特性并且使用所述道路特性来确定所述更新的目标速度。

IPR2025-00943
Tesla EX1002 Page 1249

用于模仿前车的系统和方法

[0001] 本申请是申请日为2014年12月4日，申请号为201480074153.9(PCT/IB2014/003151)，发明名称为“用于模仿前车的系统和方法”的专利申请的分案申请。

技术领域

[0002] 所述本公开总体涉及自主车辆导航，更具体地，涉及使用相机提供自主车辆导航特征的系统和方法。

背景技术

[0003] 随着技术的不断进步，能够在道路上导航的完全自主车辆的目标是即将达到。首先，自主车辆能够识别其环境并且在无操作人员输入的情况下导航。自主车辆还可考虑各种因素并且基于这些因素做出适当的决定，以安全而准确地到达预期目的地。例如，当车辆惯常地在道路上行进时会遇到诸如其它车辆和行人的各种对象。自主驱动系统可在车辆环境中辨识这些对象，并且采取适当的而及时的动作来避免碰撞。另外，自主驱动系统可识别其它指示器，诸如交通信号、交通标志和车道标记，所述其它指示器规范车辆移动(例如，当车辆必须停止以及可行驶时，车辆不能超过的速度，车辆必须定位在道路上，等等)。自主驱动系统可能需要确定车辆何时应该变道、路口转弯、改变道路，等等。根据这些示例显而易见的是，许多因素可能需要解决，以便提供能够安全而准确地导航的自主车辆。

发明内容

[0004] 与本公开相一致的实施方式提供用于自主车辆导航的系统和方法。公开的实施方式可使用相机来提供自主车辆导航特征。例如，与公开的实施方式相一致，公开的系统可包括一个、两个或更多个相机，用于监测车辆环境并且基于对由所述相机中的一个或多个捕获的图像的分析来产生导航响应。

[0005] 与公开的实施方式相一致，提供了一种用于车辆的驾驶员辅助导航系统。所述系统可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像确定在所述车辆的第一侧上的第一车道约束；根据所述多个图像确定与在所述车辆的所述第一侧相反的所述车辆的第二侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述车辆行进的车道，并且其中第一距离对应于所述车辆的所述第一侧与所述第一车道约束之间的距离，并且第二距离对应于所述车辆的所述第二侧与所述第二车道约束之间的距离；基于所述多个图像来确定车道偏移条件是否存在于所述车辆的所述第一侧上；如果车道偏移条件存在于所述车辆的所述第一侧上，则致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离大于所述第二距离；基于所述多个图像来确定车道偏移条件是否存在于所述车辆的所述第二侧上；并且如果车道偏移条件存在于所述车辆的所述第二侧上，则致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离小于所述第二距离。

IPR2025-00943
Tesla EX1002 Page 1250

[0006] 与公开的实施方式相一致，所述至少一个处理装置可进一步构造成如果基于所述多个图像确定车道偏移条件存在于所述车辆的所述第一侧和第二侧两者上，则致使所述第一距离与所述第二距离大致相同。所述至少一个处理装置可构造成基于所述多个图像来确定车道偏移条件是否基本上不存在，并且如果所述车道偏移条件基本上不存在，则致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离基本上等于所述第二距离。如果所述驾驶员辅助导航系统确定车道偏移条件存在，则所述至少一个处理装置可构造成致使所述第一距离比所述第二距离大至少1.25倍。替代地，所述至少一个处理装置可进一步构造成致使所述第一距离比所述第二距离大至少2倍。如果所述驾驶员辅助导航系统确定所述车道偏移条件包括弯路，则所述至少一个处理装置可构造成致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离朝向弯道的相对于所述车辆的外侧，并且所述第二距离朝向所述弯道的相对于所述车辆的内侧。如果所述驾驶员辅助导航系统确定所述车道偏移条件包括在所述车道的一侧上存在一个或多个对象，则所述至少一个处理装置可构造成致使所述车辆在所述第一车道约束和第二车道约束内行进，使得朝向所述一个或多个对象提供大于所述第二距离的所述第一距离。所述一个或多个对象可包括停放的车、墙和/或行人，并且可具有距路面大于10cm的高度。所述第一车道约束和第二车道约束中的至少一个可包括标记在路面上的线、路面的路边或者路面宽度的确定的中点。确定车道偏移条件是否存在于所述车辆的所述第一侧可基于目标对象在所述车辆的所述第一侧上的位置来执行。致使所述车辆在所述第一车道约束和第二车道约束内行进使得所述第一距离大于所述第二距离可基于所述车辆的驾驶员的输入而撤销。

[0007] 与另一公开的实施方式相一致，一种车辆可包括：第一车辆侧；与所述第一车辆侧相反的第二车辆侧；至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像确定在所述第一车辆侧上的第一车道约束；根据所述多个图像确定在所述第二车辆侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述车辆行进的车道，并且其中第一距离对应于所述第一车辆侧与所述第一车道约束之间的距离，并且第二距离对应于所述第二车辆侧与所述第二车道约束之间的距离；基于所述多个图像来确定车道偏移条件是否存在于所述第一车辆侧上；如果车道偏移条件存在于所述第一车辆侧上，则致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离大于所述第二距离；基于所述多个图像来确定车道偏移条件是否存在于所述第二车辆侧上；并且如果车道偏移条件存在于所述第二车辆侧上，则致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离小于所述第二距离。

[0008] 与公开的实施方式相一致，所述至少一个处理装置可进一步构造成如果基于所述多个图像确定车道偏移条件存在于所述第一车辆侧和第二车辆侧两者上，则致使所述第一距离与所述第二距离大致相同。所述至少一个处理装置可构造成基于所述多个图像来确定车道偏移条件是否基本上不存在，并且如果所述车道偏移条件基本上不存在，则致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离基本上等于所述第二距离。所述车道偏移条件可包括弯路，并且所述至少一个处理装置可构造成致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离朝向弯道的相对于所述车

IPR2025-00943
Tesla EX1002 Page 1251

辆的外侧，并且所述第二距离朝向所述弯道的相对于所述车辆的内侧。所述车道偏移条件可包括在所述车道的一侧上存在一个或多个对象，并且所述至少一个处理装置可构造成致使所述车辆在所述第一车道约束和第二车道约束内行进，使得朝向所述一个或多个对象提供大于所述第二距离的所述第一距离。所述第一车道约束和第二车道约束中的至少一个可包括路面的路边和/或路面宽度的确定的中点。确定车道偏移条件是否存在于所述车辆的第一侧可基于目标对象在所述车辆的第一侧上的位置来执行。致使所述车辆在所述第一车道约束和第二车道约束内行进使得所述第一距离大于所述第二距离可基于所述车辆的驾驶员的输入而撤销。

[0009] 与另一公开的实施方式相一致，提供了一种用于导航车辆的方法。所述方法可包括：使用至少一个图像捕获装置来获取车辆附近区域的多个图像；根据所述多个图像确定所述车辆的第一侧上的第一车道约束；根据所述多个图像确定与所述车辆的所述第一侧相反的所述车辆的第二侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述车辆行进的车道，并且其中第一距离对应于所述车辆的所述第一侧与所述第一车道约束之间的距离，并且第二距离对应于所述车辆的所述第二侧与所述第二车道约束之间的距离；基于所述多个图像来确定车道偏移条件是否存在于所述车辆的所述第一侧上；如果车道偏移条件存在于所述车辆的所述第一侧上，则致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离大于所述第二距离；基于所述多个图像来确定车道偏移条件是否存在于所述车辆的所述第二侧上；以及如果车道偏移条件存在于所述车辆的所述第二侧上，则致使所述车辆在所述第一车道约束和第二车道约束内行进，使得所述第一距离小于所述第二距离。

[0010] 与公开的实施方式相一致，提供了一种用于车辆的驾驶员辅助导航系统。所述系统可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像确定多个可用行进车道之中的当前行进车道；以及如果所述当前行进车道与预定的默认行进车道不相同则致使所述车辆变道。

[0011] 与公开的实施方式相一致，所述预定的默认行进车道可以是所述多个可用行进车道之中的最右车道。所述至少一个处理装置可构造成在致使所述车辆变道之前激活转向信号。所述至少一个处理装置可构造成在致使所述车辆变道之前产生语音通知。所述至少一个处理装置可构造成在致使所述车辆变道之前产生语音指示。可基于所述车辆附近区域的所述多个图像、基于从全球定位系统接收的数据和/或基于从用户接收的输入来确定所述预定的默认行进车道。

[0012] 与另一公开的实施方式相一致，一种车辆可包括：主体；至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像确定多个可用行进车道之中的当前行进车道；以及如果所述当前行进车道与预定的默认行进车道不相同则致使所述车辆变道。

[0013] 与公开的实施方式相一致，所述预定的默认行进车道可以是所述多个可用行进车道之中的最右车道。所述至少一个处理装置可构造成在致使所述车辆变道之前激活转向信号。所述至少一个处理装置可构造成在致使所述车辆变道之前产生语音通知。所述至少一

IPR2025-00943
Tesla EX1002 Page 1252

个处理装置可构造成在致使所述车辆变道之前产生语音指示。可基于所述车辆附近区域的所述多个图像、基于从全球定位系统接收的数据和/或基于从用户接收的输入来确定所述预定的默认行进车道。

[0014] 与另一公开的实施方式相一致，提供了一种用于导航车辆的方法。所述方法可包括：使用至少一个图像捕获装置来获取车辆附近区域的多个图像；根据所述多个图像确定多个可用行进车道之中的当前行进车道；以及如果所述当前行进车道与预定的默认行进车道不相同则致使所述车辆变道。

[0015] 与公开的实施方式相一致，所述预定的默认行进车道可以是所述多个可用行进车道之中的最右车道。可在致使所述车辆变道之前激活转向信号。可基于所述车辆附近区域的所述多个图像来确定所述预定的默认行进车道。

[0016] 与公开的实施方式相一致，提供了一种用于车辆的驾驶员辅助导航系统。所述系统可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；基于地图数据和车辆位置信息来辨识待导航的弯道；基于反映在所述地图数据中的弯道的至少一个特性来确定所述车辆的初始目标速度；将所述车辆的速度调整至所述初始目标速度；基于所述多个图像来确定弯道的观察到的一个或多个特性；基于所述弯道的观察到的一个或多个特性来确定更新的目标速度；以及将所述车辆的速度调整至所述更新的目标速度。

[0017] 与公开的实施方式相一致，所述弯道的至少一个特性可包括根据所述地图数据确定的所述弯道的预期曲率半径。所述初始目标速度可基于所述弯道的所述预期曲率半径和侧向加速极限来确定。所述侧向加速极限可介于约0.2米/秒^{2}到3.0米/秒^{2}之间。替代地，所述侧向加速极限可介于约1.5米/秒^{2}到3.0米/秒^{2}之间。所述侧向加速极限可基于由所述多个图像确定的道路宽度来调整。所述弯道的观察到的一个或多个特性可包括通过分析所述多个图像确定的所述弯道的实际曲率半径。所述更新的目标速度可基于所述弯道的所述实际曲率半径和侧向加速极限来确定。所述弯道的观察到的一个或多个特性可包括通过分析所述多个图像确定的所述弯道的实际曲率半径的变化率。所述弯道的观察到的一个或多个特性可包括通过分析所述多个图像确定的与所述弯道关联的转弯倾斜度。所述弯道的观察到的一个或多个特性可包括通过分析所述多个图像确定的与由一个或多个道路标志传达的所述弯道相关的信息。所述初始目标速度还可基于所述至少一个图像捕获装置的焦距范围。所述初始目标速度还可基于确定交通灯是否沿着所述弯道定位。所述初始目标速度还可基于确定交通灯是否位于所述弯道之后的预定距离内。

[0018] 与另一公开的实施方式相一致，一种车辆可包括：主体；至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；基于地图数据和车辆位置信息来辨识待导航的弯道；基于反映在所述地图数据中的弯道的至少一个特性来确定所述车辆的初始目标速度；将所述车辆的速度调整至所述初始目标速度；基于所述多个图像来确定弯道的观察到的一个或多个特性；基于所述弯道的观察到的一个或多个特性来确定更新的目标速度；以及将所述车辆的速度调整至所述更新的目标速度。

IPR2025-00943
Tesla EX1002 Page 1253

[0019] 与公开的实施方式相一致，所述弯道的至少一个特性可包括根据所述地图数据确定的所述弯道的预期曲率半径。所述初始目标速度可基于所述弯道的所述预期曲率半径和侧向加速极限来确定。所述侧向加速极限可基于由所述多个图像确定的道路宽度来调整。所述弯道的观察到的一个或多个特性可包括通过分析所述多个图像确定的所述弯道的实际曲率半径。所述更新的目标速度可基于所述弯道的所述实际曲率半径和侧向加速极限来确定。所述弯道的观察到的一个或多个特性可包括通过分析所述多个图像确定的所述弯道的实际曲率半径的变化率。所述弯道的观察到的一个或多个特性可包括通过分析所述多个图像确定的与所述弯道关联的转弯倾斜度。所述弯道的观察到的一个或多个特性可包括通过分析所述多个图像确定的与由一个或多个道路标志传达的所述弯道相关的信息。所述初始目标速度还可基于所述至少一个图像捕获装置的焦距范围。所述初始目标速度还可基于确定交通灯是否沿着所述弯道定位。所述初始目标速度还可基于确定交通灯是否位于所述弯道之后的预定距离内。

[0020] 与另一公开的实施方式相一致，提供了一种用于导航车辆的方法。所述方法可包括：使用至少一个图像捕获装置来获取车辆附近区域的多个图像；基于地图数据和车辆位置信息来辨识待导航的弯道；基于反映在所述地图数据中的弯道的至少一个特性来确定所述车辆的初始目标速度；将所述车辆的速度调整至所述初始目标速度；基于所述多个图像来确定弯道的观察到的一个或多个特性；基于所述弯道的观察到的一个或多个特性来确定更新的目标速度；以及将所述车辆的速度调整至所述更新的目标速度。

[0021] 与公开的实施方式相一致，提供了一种用于主车辆的驾驶员辅助导航系统。所述系统可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取主车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像中的至少一些图像来确定所述主车辆的第一侧上的第一车道约束；根据所述多个图像中的至少一些图像来确定与所述主车辆的所述第一侧相反的所述主车辆的第二侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述主车辆行进的车道；致使所述主车辆在所述第一车道约束和第二车道约束内行进；在所述多个图像中定位前车；基于所述多个图像来确定所述前车采取的至少一个动作；以及致使所述主车辆模仿所述前车的所述至少一个动作。

[0022] 与公开的实施方式相一致，所述第一车道约束和第二车道约束中的至少一个可包括标记在路面上的线。所述第一车道约束和第二车道约束中的至少一个可包括路面的路边。所述第一车道约束和第二车道约束中的至少一个可包括路面宽度的确定的中点。所述前车采取的所述至少一个动作可包括转弯，并且所述至少一个处理装置可致使所述主车辆执行开始于所述前车启动其转弯的地点处的类似转弯。所述前车采取的所述至少一个动作可包括加速。所述前车采取的所述至少一个动作可包括制动。所述前车采取的所述至少一个动作可包括在所述前车的邻近所述第一车道约束的那一侧提供第一偏移距离，所述第一偏移距离不同于所述前车的邻近所述第二车道约束的那一侧的第二偏移距离。所述至少一个处理装置可进一步构造成基于所述多个图像或者基于地图数据和位置信息来确定所述前车是否在路口处转弯，并且如果确定所述前车是在路口处转弯，则拒绝模仿所述前车的所述转弯。所述至少一个处理装置可构造成基于所述多个图像或者基于地图数据和位置信息来确定所述前车是否在路口处转弯、确定所述前车是否在所述路口内变道，如果确定所

IPR2025-00943
Tesla EX1002 Page 1254

述前车是在路口处转弯并且在所述路口内不变道则模仿所述前车的所述转弯，并且如果确定所述前车是在路口处转弯但在所述路口内变道则拒绝模仿所述前车的所述转弯。所述至少一个处理装置可构造成：基于所述多个图像来获取第一前车，基于所述多个图像来获取第二前车，并且如果基于所述多个图像确定所述第一前车继续行进的路径的转弯半径小于所述第二前车跟随的路径，则致使所述主车辆模仿所述第一前车。所述至少一个处理装置可构造成：确定所述前车的位置信息并且将所述位置信息与预定地图数据比较，通过跟踪所述前车相对于所述预定地图数据的移动为所述前车创建导航历史，并且基于所述前车的所述导航历史来确定是否致使所述主车辆模仿所述前车的所述至少一个动作。所述至少一个处理装置可构造成：当所述导航历史指示所述前车已成功导航至少一个前转弯时，甚至在基于所述多个图像未检测到转弯的情况下，致使所述主车辆模仿所述前车的转弯。

[0023] 与另一公开的实施方式相一致，一种主车辆可包括：主体；至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取主车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像中的至少一些图像来确定在所述主车辆的第一侧上的第一车道约束；根据所述多个图像中的至少一些图像来确定在与所述主车辆的所述第一侧相反的所述主车辆的第二侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述主车辆行进的车道；致使所述主车辆在所述第一车道约束和第二车道约束内行进；在所述多个图像中定位前车；基于所述多个图像来确定所述前车采取的至少一个动作；以及致使所述主车辆模仿所述前车的所述至少一个动作。

[0024] 与公开的实施方式相一致，所述第一车道约束和第二车道约束中的至少一个可包括标记在路面上的线。所述第一车道约束和第二车道约束中的至少一个可包括路面的路边。所述第一车道约束和第二车道约束中的至少一个可包括路面宽度的确定的中点。所述前车采取的所述至少一个动作可包括转弯，并且所述至少一个处理装置可致使所述主车辆执行开始于所述前车启动其转弯的地点处的类似转弯。所述前车采取的所述至少一个动作可包括加速。所述前车采取的所述至少一个动作可包括制动。所述前车采取的所述至少一个动作可包括在所述前车的邻近所述第一车道约束的那一侧提供第一偏移距离，所述第一偏移距离不同于所述前车的邻近所述第二车道约束的那一侧的第二偏移距离。所述至少一个处理装置可进一步构造成基于所述多个图像或者基于地图数据和位置信息来确定所述前车是否在路口处转弯，并且如果确定所述前车是在路口处转弯，则拒绝模仿所述前车的所述转弯。所述至少一个处理装置可构造成基于所述多个图像或者基于地图数据和位置信息来确定所述前车是否在路口处转弯、确定所述前车是否在所述路口内变道，如果确定所述前车是在路口处转弯并且在所述路口内不变道则模仿所述前车的所述转弯，并且如果确定所述前车是在路口处转弯但在所述路口内变道则拒绝模仿所述前车的所述转弯。所述至少一个处理装置可构造成：基于所述多个图像来获取第一前车，基于所述多个图像来获取第二前车，并且如果基于所述多个图像确定所述第一前车继续行进的路径的转弯半径小于所述第二前车跟随的路径，则致使所述主车辆模仿所述第一前车。所述至少一个处理装置可构造成：确定所述前车的位置信息并且将所述位置信息与预定地图数据比较，通过跟踪所述前车相对于所述预定地图数据的移动为所述前车创建导航历史，并且基于所述前车的所述导航历史来确定是否致使所述主车辆模仿所述前车的所述至少一个动作。所述至少一

IPR2025-00943
Tesla EX1002 Page 1255

个处理装置可构造成：当所述导航历史指示所述前车已成功导航至少一个前转弯时，甚至在基于所述多个图像未检测到转弯的情况下，致使所述主车辆模仿所述前车的转弯。

[0025] 与另一公开的实施方式相一致，提供了一种用于导航主车辆的方法。所述方法可包括：使用至少一个图像捕获装置来获取主车辆附近区域的多个图像；根据所述多个图像中的至少一些图像来确定在所述主车辆的第一侧上的第一车道约束；根据所述多个图像中的至少一些图像来确定在与所述主车辆的所述第一侧相反的所述主车辆的第二侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述主车辆行进的车道；致使所述主车辆在所述第一车道约束和第二车道约束内行进；在所述多个图像中定位前车；基于所述多个图像来确定所述前车采取的至少一个动作；以及致使所述主车辆模仿所述前车的所述至少一个动作。

[0026] 与公开的实施方式相一致，提供了一种用于用户车辆的驾驶员辅助导航系统。所述系统可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取用户车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像确定在所述用户车辆的第一侧上的第一车道约束；根据所述多个图像确定在与所述用户车辆的所述第一侧相反的所述用户车辆的第二侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述用户车辆行进的车道；基于所述多个图像来获取目标车辆；如果确定所述目标车辆行进的车道不同于所述用户车辆行进的车道则使所述用户车辆能够超越所述目标车辆；基于所述多个图像来监测所述目标车辆的位置；以及如果确定所述目标车辆正进入所述用户车辆行进的车道，则致使所述用户车辆在完成所述超越之前中止所述超越。

[0027] 与公开的实施方式相一致，致使所述用户车辆中止所述超越的步骤可包括致使施加所述用户车辆中的制动器。致使所述用户车辆中止所述超越的步骤可包括产生语音通知。致使所述用户车辆中止所述超越的步骤可包括致使所述用户车辆停止加速。监测所述目标车辆的位置的步骤可包括估算所述目标车辆的速度。所述第一车道约束和第二车道约束中的至少一个可包括标记在路面上的线、路面的路边和/或路面宽度的确定的中点。

[0028] 与另一公开的实施方式相一致，一种用户车辆可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取用户车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像确定在所述用户车辆的第一侧上的第一车道约束；根据所述多个图像确定在与所述用户车辆的所述第一侧相对的所述用户车辆的第二侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述用户车辆行进的车道；基于所述多个图像来获取目标车辆；如果确定所述目标车辆行进的车道不同于所述用户车辆行进的车道则使所述用户车辆能够超越所述目标车辆；基于所述多个图像来监测所述目标车辆的位置；以及如果确定所述目标车辆正进入所述用户车辆行进的车道，则致使所述用户车辆在完成所述超越之前中止所述超越。

[0029] 与公开的实施方式相一致，致使所述用户车辆中止所述超越可包括致使施加所述用户车辆中的制动器。致使所述用户车辆中止所述超越可包括产生语音通知。致使所述用户车辆中止所述超越可包括致使所述用户车辆停止加速。监测所述目标车辆的位置可包括估算所述目标车辆的速度。所述第一车道约束和第二车道约束中的至少一个可包括标记在

IPR2025-00943
Tesla EX1002 Page 1256

路面上的线、路面的路边和/或路面宽度的确定的中点。

[0030] 与另一公开的实施方式相一致，提供了一种用于导航用户车辆的方法。所述方法可包括：使用至少一个图像捕获装置来获取用户车辆附近区域的多个图像；根据所述多个图像确定在所述用户车辆的第一侧上的第一车道约束；根据所述多个图像确定与在所述用户车辆的所述第一侧相反的所述用户车辆的第二侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述用户车辆行进的车道；基于所述多个图像来获取目标车辆；如果确定所述目标车辆行进的车道不同于所述用户车辆行进的车道则使所述用户车辆能够超越所述目标车辆；基于所述多个图像来监测所述目标车辆的位置；以及如果确定所述目标车辆正进入所述用户车辆行进的车道，则致使所述用户车辆在完成所述超越之前中止所述超越。

[0031] 与公开的实施方式相一致，致使所述用户车辆中止所述超越的步骤可包括致使施加所述用户车辆中的制动器。致使所述用户车辆中止所述超越的步骤可包括致使所述用户车辆停止加速。所述第一车道约束和第二车道约束中的至少一个可包括标记在路面上的线、路面的路边和/或路面宽度的确定的中点。

[0032] 与公开的实施方式相一致，提供了一种用于用户车辆的驾驶员辅助导航系统。所述系统可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取用户车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像确定在所述用户车辆的第一侧上的第一车道约束；根据所述多个图像确定在与所述用户车辆的所述第一侧相反的所述用户车辆的第二侧上的第二车道约束，其中所述第一车道约束和第二车道约束限定所述用户车辆行进的车道；基于所述多个图像来确定是否有侵入车辆从所述用户车辆的一侧接近；以及致使所述用户车辆维持当前速度并且在所述第一车道约束和第二车道约束内行进，使得在所述用户车辆的所述侵入车辆正接近的那一侧的第一距离大于第二距离。

[0033] 与公开的实施方式相一致，所述至少一个处理装置可进一步构造成如果对所述多个图像的分析指示所述侵入车辆已越过至少一个碰撞阈值则致使所述用户车辆采取躲避动作。所述躲避动作可包括变道、变动路线和/或制动。所述至少一个碰撞阈值可包括在所述用户车辆与所述侵入车辆之间维持的最小预定距离。所述至少一个碰撞阈值可包括将要碰撞的时间。所述至少一个处理装置可构造成在完成所述躲避动作之后致使所述用户车辆返回到所述用户车辆最初行进的车道的中央。所述至少一个图像捕获装置可具有至少140度的视场。所述至少一个处理装置可构造成在响应于侵入车辆进行制动之后致使所述用户车辆恢复预设速度。所述至少一个处理装置可构造成致使所述用户车辆返回到所述第一车道约束和第二车道约束内的车道地点，使得在响应于侵入车辆而变动路线之后，所述用户车辆与所述第一车道约束之间的所述第一距离基本上等于所述用户车辆与所述第二车道约束之间的所述第二距离。

[0034] 与另一公开的实施方式相一致，一种用户车辆可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取用户车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；根据所述多个图像确定在所述用户车辆的第一侧上的第一车道约束；根据所述多个图像确定在与所述用户车辆的所述第一侧相反的所述用户车辆的第二侧上的第二车道约束，其中

IPR2025-00943
Tesla EX1002 Page 1257

所述第一车道约束和第二车道约束限定所述用户车辆行进的车道;基于所述多个图像来确定是否有侵入车辆从所述用户车辆的一侧接近;以及致使所述用户车辆维持当前速度并且在所述第一车道约束和第二车道约束内行进,使得在所述用户车辆的所述侵入车辆正接近的那一侧的第一距离大于第二距离。

[0035] 与公开的实施方式相一致,所述至少一个处理装置可进一步构造成如果所述对多个图像的分析指示所述侵入车辆已越过至少一个碰撞阈值则致使所述用户车辆采取躲避动作。所述躲避动作可包括变道、变动路线和/或制动。所述至少一个碰撞阈值可包括在所述用户车辆与所述侵入车辆之间维持的最小预定距离。所述至少一个碰撞阈值可包括将要碰撞的时间。所述至少一个处理装置可构造成在完成所述躲避动作之后致使所述用户车辆返回到所述用户车辆最初行进的车道的中央。所述至少一个图像捕获装置可具有至少140度的视场。所述至少一个处理装置可构造成在响应于侵入车辆进行制动之后致使所述用户车辆恢复预设速度。所述至少一个处理装置可构造成致使所述用户车辆返回到所述第一车道约束和第二车道约束内的车道地点,使得在响应于侵入车辆而变动路线之后,所述用户车辆与所述第一车道约束之间的所述第一距离基本上等于所述用户车辆与所述第二车道约束之间的所述第二距离。

[0036] 与另一公开的实施方式相一致,提供了一种用于导航用户车辆的方法。所述方法可包括:使用至少一个图像捕获装置来获取用户车辆附近区域的多个图像;根据所述多个图像确定在所述用户车辆的第一侧上的第一车道约束;根据所述多个图像确定在与所述用户车辆的所述第一侧相反的所述用户车辆的第二侧上的第二车道约束,其中所述第一车道约束和第二车道约束限定所述用户车辆行进的车道;基于所述多个图像来确定是否有侵入车辆从所述用户车辆的一侧接近;以及致使所述用户车辆维持当前速度并且在所述第一车道约束和第二车道约束内行进,使得在所述用户车辆的所述侵入车辆正接近的那一侧的第一距离大于第二距离。

[0037] 与公开的实施方式相一致,提供了一种用于主车辆的驾驶员辅助导航系统。所述系统可包括:至少一个图像捕获装置,所述至少一个图像捕获装置被构造成获取主车辆附近区域的多个图像;数据接口;以及至少一个处理装置。所述至少一个处理装置可构造成:经由所述数据接口接收所述多个图像;识别所述多个图像内的目标对象;经由所述多个图像监测所述目标对象的运动以及所述主车辆与所述目标对象之间的距离;基于监测的运动以及所述主车辆与所述目标对象之间的距离来确定所述主车辆与所述目标对象之间的拦截时间的指示器;以及基于所述拦截时间与多个预定拦截阈值的比较来产生所述主车辆的响应。

[0038] 与公开的实施方式相一致,所述目标对象可以是在与所述主车辆不同的车道中行进的目标车辆。所述拦截时间可基于由所述多个图像确定的被跟踪对象的速度的指示器。所述拦截时间可基于由所述多个图像确定的所述被跟踪对象的加速度的指示器。所述至少一个处理装置可进一步构造成确定所述目标对象是否在与所述主车辆相同的车道中行进,并且如果确定所述目标对象是在与所述主车辆不同的车道中行进则放弃所述主车辆的所述响应。确定所述目标对象是否在与所述主车辆相同的车道中行进可基于所述多个图像。确定所述目标对象是否在与所述主车辆相同的车道中行进可基于所述主车辆的预定地图数据或偏航率中的至少一者。如果拦截时间的所述指示器指示所述目标对象在相对于所述

IPR2025-00943
Tesla EX1002 Page 1258

车辆的第一拦截阈值的外侧，则所述响应可基于与所述目标对象关联的平均运动轮廓。如果拦截时间的所述指示器指示所述目标对象在相对于所述车辆的所述第一拦截阈值与第二拦截阈值之间，则所述响应可基于所述目标对象的实际运动。如果拦截时间的所述指示器指示所述目标对象在所述第二拦截阈值内，则所述响应可包括紧急避让动作。所述紧急避让动作可包括变道和/或紧急制动。

[0039] 与另一公开的实施方式相一致，一种主车辆可包括：主体；至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取主车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；识别所述多个图像内的目标对象；经由所述多个图像监测所述目标对象的运动以及所述主车辆与所述目标对象之间的距离；基于监测的运动以及所述主车辆与所述目标对象之间的距离来确定所述主车辆与所述目标对象之间的拦截时间的指示器；以及基于所述拦截时间与多个预定拦截阈值的比较来产生所述主车辆的响应。

[0040] 与公开的实施方式相一致，所述目标对象可以是在与所述主车辆不同的车道中行进的目标车辆。所述拦截时间可基于由所述多个图像确定的被跟踪对象的速度的指示器。所述拦截时间可基于由所述多个图像确定的所述被跟踪对象的加速度的指示器。所述至少一个处理装置可进一步构造成确定所述目标对象是否在与所述主车辆相同的车道中行进，并且如果确定所述目标对象是在与所述主车辆不同的车道中行进则放弃所述主车辆的所述响应。确定所述目标对象是否在与所述主车辆相同的车道中行进可基于所述多个图像。确定所述目标对象是否在与所述主车辆相同的车道中行进可基于所述主车辆的预定地图数据或偏航率中的至少一者。如果拦截时间的所述指示器指示所述目标对象在相对于所述车辆的第一拦截阈值的外侧，则所述响应可基于与所述目标对象关联的平均运动轮廓。如果拦截时间的所述指示器指示所述目标对象在相对于所述车辆的所述第一拦截阈值与第二拦截阈值之间，则所述响应可基于所述目标对象的实际运动。如果拦截时间的所述指示器指示所述目标对象在所述第二拦截阈值内，则所述响应可包括紧急避让动作。所述紧急避让动作可包括变道和/或紧急制动。

[0041] 与另一公开的实施方式相一致，一种用于导航主车辆的方法可包括：经由至少一个图像捕获装置获取主车辆附近区域的多个图像；识别所述多个图像内的目标对象；基于所述多个图像来监测所述目标对象的运动以及所述主车辆与所述目标对象之间的距离；基于监测的运动以及所述主车辆与所述目标对象之间的距离来确定所述主车辆与所述目标对象之间的拦截时间的指示器；以及基于所述拦截时间与多个预定拦截阈值的比较来产生所述主车辆的响应。

[0042] 与公开的实施方式相一致，提供了一种用于车辆的驾驶员辅助导航系统。所述系统可包括：至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；基于所述多个图像的分析来识别用于停止所述车辆的触发器；以及基于识别的所述触发器，根据制动轮廓来致使所述车辆停止，如基于对所述多个图像的分析确定的，所述制动轮廓包括与第一减速率关联的第一区段、具有小于所述第一减速率的第二减速率的第二区段以及在接近目标停止地点时使制动水平降低的第三区段。

[0043] 与公开的实施方式相一致，到达所述目标停止地点可伴随着制动压力的释放。所

IPR2025-00943
Tesla EX1002 Page 1259

述触发器可包括对象。所述对象可包括所述车辆前方的另一车辆。所述触发器可包括交通灯信号、停止标志和/或限速标志。所述至少一个处理装置可进一步构造成提供识别所述触发器的语音通知。所述至少一个处理装置可构造成产生指示所述车辆正制动的语音通知。

[0044] 与另一公开的实施方式相一致，一种车辆可包括：主体；至少一个图像捕获装置，所述至少一个图像捕获装置被构造成获取车辆附近区域的多个图像；数据接口；以及至少一个处理装置。所述至少一个处理装置可构造成：经由所述数据接口接收所述多个图像；基于对所述多个图像的分析来识别用于停止所述车辆的触发器；以及基于识别的所述触发器，根据制动轮廓来致使所述车辆停止，如基于所述多个图像的分析确定的，所述制动轮廓包括与第一减速率关联的第一区段、具有小于所述第一减速率的第二减速率的第二区段以及在接近目标停止地点时使制动水平降低的第三区段。

[0045] 与公开的实施方式相一致，到达所述目标停止地点可伴随着制动压力的释放。所述触发器可包括对象。所述对象可包括所述车辆前方的另一车辆。所述触发器可包括交通灯信号、停止标志和/或限速标志。所述至少一个处理装置可进一步构造成提供识别所述触发器的语音通知。所述至少一个处理装置可构造成产生指示所述车辆正制动的语音通知。

[0046] 与另一公开的实施方式相一致，提供了一种用于导航车辆的方法。所述方法可包括：经由至少一个图像捕获装置获取车辆附近区域的多个图像；基于对所述多个图像的分析来识别用于停止所述车辆的触发器；以及基于识别的所述触发器，根据制动轮廓来致使所述车辆停止，如基于所述多个图像的分析确定的，所述制动轮廓包括与第一减速率关联的第一区段、具有小于所述第一减速率的第二减速率的第二区段以及在接近目标停止地点时使制动水平降低的第三区段。

[0047] 与其它公开的实施方式相一致，非临时性计算机可读的存储介质可存储程序指令，该程序指令由至少一个处理装置执行并且执行本文中描述的任何方法。

[0048] 前面的总体描述和以下详细描述仅是示例性和说明性的并不限制权利要求书。

附图说明

[0049] 并入本公开并构成本公开一部分的附图图示了各种公开的实施方式。在图中：

[0050] 图1是与公开的实施方式相一致的示例性系统的概略表示。

[0051] 图2A是包括与公开的实施方式相一致的系统的示例性车辆的概略侧视图表示。

[0052] 图2B是与公开的实施方式相一致的图2A所示车辆和系统的概略顶视图表示。

[0053] 图2C是包括与公开的实施方式相一致的系统的车辆的另一实施方式的概略顶视图表示。

[0054] 图2D是包括与公开的实施方式相一致的系统的车辆的又一实施方式的概略顶视图表示。

[0055] 图2E是包括与公开的实施方式相一致的系统的车辆的又一实施方式的概略顶视图表示。

[0056] 图2F是与公开的实施方式相一致的示例性车辆控制系统的概略表示。

[0057] 图3A是与公开的实施方式相一致的车辆的内部的概略表示，包括用于车辆成像系统的后视镜和用户接口。

[0058] 图3B是与公开的实施方式相一致的构造成定位在后视镜后面并且抵着车辆挡风

IPR2025-00943
Tesla EX1002 Page 1260

玻璃的相机安装件的示例的图示。

[0059] 图3C是与公开的实施方式相一致的图3B所示的视角不同的相机安装件的图示。

[0060] 图3D是与公开的实施方式相一致的构造成定位在后视镜后面并且抵着车辆挡风玻璃的相机安装件的示例的图示。

[0061] 图4是与公开的实施方式相一致的构造成存储用于执行一个或多个操作的指令的存储器的示例性框图。

[0062] 图5A是示出与公开的实施方式相一致的用于基于单眼图像分析来产生一个或多个导航响应的示例性过程的流程图。

[0063] 图5B是示出与公开的实施方式相一致的用于检测一组图像中的一个或多个车辆和/或行人的示例性过程的流程图。

[0064] 图5C是示出与公开的实施方式相一致的用于检测一组图像中的路标和/或车道几何信息的示例性过程的流程图。

[0065] 图5D是示出与公开的实施方式相一致的用于检测一组图像中的交通灯的示例性过程的流程图。

[0066] 图5E是示出与公开的实施方式相一致的用于基于车辆路径来产生一个或多个导航响应的示例性过程的流程图。

[0067] 图5F是示出与公开的实施方式相一致的用于确定前车是否变道的示例性过程的流程图。

[0068] 图6是示出与公开的实施方式相一致的用于基于立体图像分析来产生一个或多个导航响应的示例性过程的流程图。

[0069] 图7是示出与公开的实施方式相一致的用于基于三组图像分析来产生一个或多个导航响应的示例性过程的流程图。

[0070] 图8A是与公开的实施方式相一致的在第一车道约束和第二车道约束内行进的示例性车辆的概略表示，其中车道偏移条件存在于第一车辆侧。

[0071] 图8B是与公开的实施方式相一致的在第一车道约束和第二车道约束内行进的示例性车辆的概略表示，其中车道偏移条件存在于第二车辆侧。

[0072] 图8C是与公开的实施方式相一致的在弯路上的第一车道约束和第二车道约束内行进的示例性车辆的概略表示。

[0073] 图8D是与公开的实施方式相一致的在第一车道约束和第二车道约束内行进的示例性车辆的概略表示，其中车道偏移条件存在于第一车辆侧和第二车辆侧两者上。

[0074] 图9是与公开的实施方式相一致的示例性导航系统的存储器的概略表示。

[0075] 图10是与公开的实施方式相一致的用于导航车辆的示例性方法的流程图。

[0076] 图11是与公开的实施方式相一致的在默认车道内行进包括导航系统的示例性车辆的概略表示。

[0077] 图12A是与公开的实施方式相一致的示例性导航系统的存储器的概略表示。

[0078] 图12B是与公开的实施方式相一致的用于导航车辆的示例性过程的流程图。

[0079] 图13A和图13B是与公开的实施方式相一致的接近以及导航具有一个或多个特性的弯道的示例性车辆的概略表示。

[0080] 图14是与公开的实施方式相一致的构造成存储用于执行一个或多个操作的指令

IPR2025-00943
Tesla EX1002 Page 1261

的存储器的示例性框图。

[0081] 图15是示出与公开的实施方式相一致的用于基于被检测弯道和观察的弯道特性来控制车辆速度的示例性过程的流程图。

[0082] 图16A至图16D是与公开的实施方式相一致的模仿前车一个或多个动作的主车辆的概略表示。

[0083] 图17是与公开的实施方式相一致的构造成存储用于执行一个或多个操作的指令的存储器的示例性框图。

[0084] 图18是示出与公开的实施方式相一致的用于致使主车辆模仿前车一个或多个动作的示例性过程的流程图。

[0085] 图19是示出与公开的实施方式相一致的用于致使主车辆拒绝模仿前车转弯的示例性过程的流程图。

[0086] 图20是示出与公开的实施方式相一致的用于致使主车辆模仿或拒绝前车转弯的示例性过程的流程图。

[0087] 图21是示出与公开的实施方式相一致的用于致使主车辆模仿前车一个或多个动作的另一示例性过程的流程图。

[0088] 图22是示出与公开的实施方式相一致的用于基于导航历史来致使主车辆模仿前车一个或多个动作的示例性过程的流程图。

[0089] 图23是与公开的实施方式相一致的在另一车辆行进所沿着的道路上的第一车道约束和第二车道约束内行进的示例性车辆的概略表示。

[0090] 图24A是与公开的实施方式相一致的存储用于在侵入车辆之中导航车辆的指令的存储器的概略表示。

[0091] 图24B是与公开的实施方式相一致的用于在侵入车辆之中导航车辆的示例性过程的流程图。

[0092] 图25是与公开的实施方式相一致的在另一车辆行进所沿着的道路上的第一车道约束和第二车道约束内行进的示例性车辆的概略表示。

[0093] 图26是与公开的实施方式相一致的存储用于检测以及响应于车辆交通侧向侵入的指令的存储器的概略表示。

[0094] 图27是与公开的实施方式相一致的用于检测以及响应于车辆交通侧向侵入的示例性过程的流程图。

[0095] 图28A和图28B是与公开的实施方式相一致的产生主车辆响应的概略表示。

[0096] 图29是与公开的实施方式相一致的构造成存储用于执行一个或多个操作的指令的存储器的示例性框图。

[0097] 图30是示出与公开的实施方式相一致的用于产生主车辆响应的示例性过程的流程图。

[0098] 图31是示出与公开的实施方式相一致的用于拒绝产生主车辆响应的示例性过程的流程图。

[0099] 图32A是与公开的实施方式相一致的使对象存在于车辆附近的示例性车辆的概略表示。

[0100] 图32B图示了与公开的实施方式相一致的示例性制动轮廓。

IPR2025-00943
Tesla EX1002 Page 1262

[0101] 图33是与公开的实施方式相一致的示例性导航系统的存储器的概略表示。

[0102] 图34是与公开的实施方式相一致的用于在导航车辆的同时实施制动轮廓的示例性方法的流程图。

具体实施方式

[0103] 以下详细描述参照附图。在可能的情况下，相同的附图标记在图中和以下描述中用来指相同或类似的部分。虽然本文中描述了若干说明性实施方式，但是可以进行修改、改写及其它实施。例如，对图中图示的部件可做出替换、添加或修改，并且本文中描述的说明性方法可通过将步骤替换、重新排序、移除或添加至公开的方法进行修改。因此，以下详细描述并不限于公开的实施方式和示例。相反，适当范围由所附权利要求限定。

[0104] 公开的实施方式提供使用相机来提供自主导航特征的系统和方法。在各种实施方式中，所述系统可包括监测车辆的环境的一个、两个或更多个相机。在一个实施方式中，所述系统可监测并调整车辆与车道边界之间的自由空间，车辆在所述车道内行进。在另一实施方式中，所述系统可选择特定车道作为车辆在行进同时使用的默认车道。在另一实施方式中，所述系统可控制车辆在不同场景(诸如转弯时)中的速度。在又一实施方式中，所述系统可模仿前车的动作。在又一实施方式中，所述系统可监测目标车辆并且在某些条件下(例如，如果目标车辆行进的车道不同于所述车辆行进的车道)使所述车辆能够超越目标车辆。在再一实施方式中，所述系统可响应于侧向侵入车辆(诸如试图进入所述车辆行进的车道中的车辆)来产生自然感觉。

[0105] 图1是与示例性公开的实施方式相一致的系统100的框图表示。取决于特定实施的要求，系统100可包括各种部件。在一些实施方式中，系统100可包括处理单元110、图像获取单元120、位置传感器130、一个或多个存储器单元140、150、地图数据库160和用户接口170。处理单元110可包括一个或多个处理装置。在一些实施方式中，处理单元110可包括应用处理器180、图像处理器190或任何其它合适的处理装置。类似地，取决于特定应用的要求，图像获取单元120可包括任何数目的图像获取装置和部件。在一些实施方式中，图像获取单元120可包括一个或多个图像捕获装置(例如，相机)，诸如图像捕获装置122、图像捕获装置124和图像捕获装置126。系统100可还包括数据接口128，该数据接口128将处理装置110和图像获取装置120通信连接。例如，数据接口128可包括用于将由图像获取装置120获取的图像数据发送到处理单元110的任何一个或多个有线和/或无线链路。

[0106] 应用处理器180和图像处理器190两者均可包括各种类型的处理装置。例如，应用处理器180和图像处理器190中的任一者或两者可包括微处理器、预处理器(诸如图像预处理器)、图形处理器、中央处理单元(CPU)、支持电路、数字信号处理器、集成电路、存储器或适于运行应用程序以及图像处理和分析的任何其它类型的装置。在一些实施方式中，应用处理器180和/或图像处理器190可包括任何类型的单核或多核处理器、移动装置微控制器、中央处理单元，等等。可使用各种处理装置，包括例如可购自诸如Intel®、AMD®等等制造商的处理器，并且可包括各种结构(例如，x86处理器、ARM®，等等)。

[0107] 在一些实施方式中，应用处理器180和/或图像处理器190可包括任何可购自Mobileye®的EyeQ系列处理器芯片。这些处理器设计各包括具有本地存储器和指令集的多处理单元。这样的处理器可包括用于接收来自多个图像传感器的图像数据的视频输入，并

IPR2025-00943
Tesla EX1002 Page 1263

且还可包括视频输出功能。在一个示例中，EyeQ2®使用在332Mhz下操作的90纳米微米技术。EyeQ2®结构组成如下：两个浮点、超线程32位RISC CPU(MIPS32® 34K®核)、五个视觉计算引擎(VCE)、三个矢量微码处理器(VMP®)、Denali 64位移动DDR控制器、128位内部超音速互连、双16位视频输入控制器和18位视频输出控制器、16通道DMA以及若干外设。MIPS34K CPU管理五个VCE、三个VMP™和DMA、第二MIPS34K CPU和多通道DMA以及其它外设。五个VCE、三个VMP®和MIPS34K CPU可以执行多功能捆绑应用程序所需要的密集的视觉计算。在另一示例中，为第三代处理器并且比EyeQ2®强大六倍的EyeQ3®可在公开的实施方式中使用。

[0108] 虽然图1描绘了包括在处理单元110中的两个分离处理装置，但是可使用或多或少的处理装置。例如，在一些实施方式中，单个处理装置可用于完成应用处理器180和图像处理器190的任务。在其它实施方式中，这些任务可由多于两个处理装置执行。

[0109] 处理单元110可包括各种类型的装置。例如，处理单元110可包括各种装置，诸如控制器、图像预处理器、中央处理单元(CPU)、支持电路、数字信号处理器、集成电路、存储器或任何其它类型的用于图像处理和分析的装置。图像预处理器可包括用于捕获、数字化并处理来自图像传感器的图像的视频处理器。CPU可包括任何数目的微控制器或微处理器。支持电路可以是本领域普遍公知的任何数目的电路，包括高速缓存、电源、时钟和输入输出电路。存储器可存储软件，该软件当由处理器执行时会控制系统的操作。存储器可包括数据库和图像处理软件。存储器可包括任何数目的随机存取存储器、只读存储器、闪速存储器、盘驱动器、光学存储介质、磁带存储介质、可去除存储介质及其它类型的存储介质。在一个实例中，存储器可从处理单元110分离。在另一实例中，存储器可集成到处理单元110中。

[0110] 每个存储器140、150可包括软件指令，该软件指令当由处理器(例如，应用处理器180和/或图像处理器190)执行时可控制系统100的各方面的操作。这些存储器单元可包括各种数据库和图像处理软件。存储器单元可包括随机存取存储器、只读存储器、闪速存储器、盘驱动器、光学存储介质、磁带存储介质、可去除存储介质和/或任何其它类型的存储介质。在一些实施方式中，存储器单元140、150可从应用处理器180和/或图像处理器190分离。在其它实施方式中，这些存储器单元可集成到应用处理器180和/或图像处理器190中。

[0111] 位置传感器130可包括适于确定与系统100的至少一个部件关联的地点的任何类型的装置。在一些实施方式中，位置传感器130可包括GPS接收器。这样的接收器可以通过处理全球定位系统卫星传播的信号来确定用户位置和速度。来自位置传感器130的位置信息可供应用处理器180和/或图像处理器190使用。

[0112] 用户接口170可包括适于向来自系统100的一个或多个用户的输入提供信息或者用于接收所述输入的任何装置。在一些实施方式中，用户接口170可包括用户输入装置，包括例如触摸屏、麦克风、键盘、指针装置、跟踪轮、相机、旋钮、按钮，等等。在这样的输入装置的情况下，用户能够通过以下手段向系统100提供信息输入或命令：键入指令或信息、提供语音命令、使用按钮、指针或眼跟踪功能在屏幕上选择菜单选项或者用于将信息与系统100通信的任何其它合适的技术。

[0113] 用户接口170可配备有一个或多个处理装置，该一个或多个处理装置被构造成向用户提供信息或从用户接收信息并且处理该信息以供例如应用处理器180使用。在一些实

IPR2025-00943
Tesla EX1002 Page 1264

施方式中，这样的处理装置可执行指令以辨识并跟踪眼移动、接收并解译语音命令、辨识并解译在触摸屏上做出的触摸和/或手势、响应于键盘输入或菜单选择，等等。在一些实施方式中，用户接口170可包括显示器、扬声器、触觉装置和/或用于向用户提供输出信息的任何其它装置。

[0114] 地图数据库160可包括用于存储对系统100有用的地图数据的任何类型的数据库。在一些实施方式中，地图数据库160可包括各种条目的在参考坐标系中与位置相关的数据，所述条目包括道路、水特征、地理特征、企业、关注点、餐厅、加油站，等等。地图数据库160不仅可存储这样的条目的地点，而且可存储与这些条目相关的描述词(包括例如，与任何存储特征关联的名称)。在一些实施方式中，地图数据库160可与系统100的其它部件物理定位。替代地或另外，地图数据库160或其一部分可相对于系统100的其它部件(例如，处理单元110)远程定位。在这样的实施方式中，来自地图数据库160的信息可通过网络的有线或无线数据连接(例如，通过蜂窝网络和/或因特网，等等)下载。

[0115] 图像捕获装置122、124和126可各包括适于从环境中捕获至少一个图像的任何类型的装置。而且，任何数目的图像捕获装置可用于获取图像以输入到图像处理器。一些实施方式可仅包括单个图像捕获装置，而其它实施方式可包括两个、三个或者甚至四个以上图像捕获装置。图像捕获装置122、124和126将在下面参照图2B至图2E进一步描述。

[0116] 系统100或其各种部件可整合到各种不同的平台中。在一些实施方式中，系统100可包括在如图2A所示的车辆200上。例如，如上面相对于图1描述的，车辆200可配备有处理单元110和系统100的任何其它部件。而在一些实施方式中，车辆200可仅配备有单个图像捕获装置(例如，相机)，在诸如结合图2B至图2E讨论的其它实施方式中，可使用多个图像捕获装置。例如，如图2A所示的车辆200的图像捕获装置122和124中的任一者可以是ADAS(先进驾驶员辅助系统)成像组的一部分。

[0117] 包括在车辆200上作为图像获取单元120的一部分的图像捕获装置可定位在任何合适的地点处。在一些实施方式中，如图2A至图2E和图3A至图3C所示，图像捕获装置122可位于后视镜附近。该位置可提供视线(类似于车辆200的驾驶员的视线)，可帮助确定对驾驶员而言什么是可见的，什么是不可见的。图像捕获装置122可定位在后视镜附近的任何地点处，但将图像捕获装置122放置在镜子的在驾驶员的那一侧上可进一步帮助获得表示驾驶员视场和/或视线的图像。

[0118] 针对图像获取单元120的图像捕获装置，还可使用其它地点。例如，图像捕获装置124可位于车辆200的保险杠之上或之中。这样的地点可尤其适合具有宽视场的图像捕获装置。位于保险杠上的图像捕获装置的视线可以不同于驾驶员的视线，因此，保险杠图像捕获装置和驾驶员可能并不总是看到相同的对象。图像捕获装置(例如，图像捕获装置122、124和126)也可位于其它地点中。例如，图像捕获装置可位于车辆200的一个或两个侧视镜之上或之中、车辆200的车顶上、车辆200的发动机罩上、车辆200的行李箱上、车辆200的侧面，安装在车辆200的任何窗口上、定位在车辆200的任何窗口的后面或前面，以及安装在车辆200的正面和/或背面的光图之中或附近，等等。

[0119] 除了图像捕获装置，车辆200可包括系统100的各种其它部件。例如，处理单元110可包括在车辆200上，与车辆的发动机控制单元(ECU)或者集成或者分离。车辆200还可配备有诸如GPS接收器的位置传感器130，并且还可包括地图数据库160以及存储器单元140和

IPR2025-00943
Tesla EX1002 Page 1265

150。

[0120] 图2A是与公开的实施方式相一致的示例性车辆成像系统的概略侧视图表示。图2B是图2A所示的实施方式的概略顶视图图示。如图2B中图示的，公开的实施方式可包括车辆200，车辆200在其主体中包括系统100，系统100具有定位在后视镜附近和/或车辆200的驾驶员附近的第一图像捕获装置122、定位在车辆200的保险杠区域(例如，保险杠区域210中的一个)之上或之中的第二图像捕获装置124，以及处理单元110。

[0121] 如图2C中图示的，图像捕获装置122和124两者均可定位在后视镜附近和/或车辆200的驾驶员附近。另外，虽然在图2B和图2C中示出了两个图像捕获装置122和124，但是应该理解的是，其它实施方式可包括多于两个图像捕获装置。例如，在图2D和图2E示出的实施方式中，第一图像捕获装置122，第二图像捕获装置124和第三图像捕获装置126被包括在车辆200的系统100中。

[0122] 如图2D中图示的，图像捕获装置122可定位在后视镜附近和/或车辆200的驾驶员附近，并且图像捕获装置124和126可定位在车辆200的保险杠区域(例如，保险杠区域210中的一个)之上或之中。并且如图2E所示，图像捕获装置122、124和126可定位在后视镜附近和/或车辆200的驾驶员座椅附近。公开的实施方式并不限于任何特定数目和构造的图像捕获装置，并且图像捕获装置可定位在车辆200内和/或上的任何适当的地点中。

[0123] 应理解的是，公开的实施方式并不限于车辆并且可以应用在其它背景中。还应理解的是，公开的实施方式并不限于特定类型的车辆200并且可应用于所有类型的车辆，包括汽车、卡车、拖车和其它类型的车辆。

[0124] 第一图像捕获装置122可包括任何合适类型的图像捕获装置。图像捕获装置122可包括光轴。在一个实例中，图像捕获装置122可包括具有全局快门的Aptina M9V024WVGA传感器。在其它实施方式中，图像捕获装置122可提供1280x960像素的分辨率并且可包括卷帘快门。图像捕获装置122可包括各种光学元件。在一些实施方式中，例如，可包括一个或多个透镜，为图像捕获装置提供期望的焦距和视场。在一些实施方式中，图像捕获装置122可与6mm透镜或12mm透镜关联。在一些实施方式中，如图2D中图示的，图像捕获装置122可构造成捕获具有期望视场(FOV)202的图像。例如，图像捕获装置122可构造成具有常规的FOV，诸如在40度到56度的范围内，包括46度FOV、50度FOV、52度FOV以上。替代地，图像捕获装置122可构造成具有23度到40度范围内的窄FOV，诸如28度FOV或36度FOV。另外，图像捕获装置122可构造成具有100度到180度范围内的宽FOV。在一些实施方式中，图像捕获装置122可包括宽角度的保险杠相机或者具有直到180度FOV。

[0125] 第一图像捕获装置122可相对于与车辆200关联的场景获取多个第一图像。多个第一图像中的每个均可被获取，作为可使用卷帘快门捕获的一系列图像扫描线。每个扫描线均可包括多个像素。

[0126] 第一图像捕获装置122可具有与获取第一系列图像扫描线中的每个关联的扫描速率。扫描速率可指这样的速率，即在所述速率下，图像传感器可以获取与包括在特定扫描线中的每个像素关联的图像数据。

[0127] 例如，图像捕获装置122、124和126可包含任何合适类型和数目的图像传感器，包括CCD传感器或CMOS传感器。在一个实施方式中，CMOS图像传感器可连同卷帘快门一起采用，使得一行中的每个像素一次读取一个，并且该行的扫描在一行接一行的基础上进行，直

IPR2025-00943
Tesla EX1002 Page 1266

到整个图像帧已被捕获。在一些实施方式中，该行可相对于帧从上到下依次捕获。

[0128] 使用卷帘快门可导致不同行中的像素在不同的时间被曝光并被捕获，这可致使被捕获的图像帧的歪斜和其它图像伪影。另一方面，当图像捕获装置122被构造成用全局或同步快门操作时，所有像素都可曝光持续相同的时间量并且在共同的曝光周期中。结果是，从采用全局快门的系统中收集的帧的图像数据表示整个FOV(诸如FOV202)在特定时间的快照。对比之下，在卷帘快门应用中，帧中的每行均被曝光，并且数据在不同的时间被捕获。由此，移动的对象可能在具有卷帘快门的图像捕获装置中呈现失真。该现象将在下面更详细地描述。

[0129] 第二图像捕获装置124和第三图像捕获装置126可以是任何类型的图像捕获装置。如同第一图像捕获装置122，图像捕获装置124和126中的每个均可包括光轴。在一个实施方式中，图像捕获装置124和126中的每个均可包括具有全局快门的Aptina M9V024 WVGA传感器。替代地，图像捕获装置124和126中的每个均可包括卷帘快门。如同图像捕获装置122，图像捕获装置124和126可构造成包括各种透镜和光学元件。在一些实施方式中，与图像捕获装置124和126关联的透镜可提供等于或窄于与图像捕获装置122关联的FOV(诸如FOV 202)的FOV(诸如FOV 204和206)。例如，图像捕获装置124和126可具有40度、30度、26度、23度、20度或以下的FOV。

[0130] 图像捕获装置124和126可相对于与车辆200的关联的场景获取多个第二图像和第三图像。多个第二图像和第三图像中的每个均可获取为可使用卷帘快门捕获的第二系列和第三系列的图像扫描线。每个扫描线或行均可具有多个像素。图像捕获装置124和126可具有与获取包括在第二系列和第三系列中的每个图像扫描线关联的第二扫描速率和第三扫描速率。

[0131] 每个图像捕获装置122、124和126均可相对于车辆200定位在任何合适的位置和取向。可选择图像捕获装置122、124和126的相对定位以帮助将从图像捕获装置获取的信息融合到一起。例如，在一些实施方式中，与图像捕获装置124关联的FOV(诸如FOV 204)可与关联图像捕获装置122的FOV(诸如FOV 202)和关联图像捕获装置126的FOV(诸如FOV 206)部分地或完全地地重叠。

[0132] 图像捕获装置122、124和126可位于车辆200上任何合适的相对高度处。在一个实例中，在图像捕获装置122、124和126之间可能存在高度差，这可提供足够的视差信息来实现立体分析。例如，如图2A所示，两个图像捕获装置122和124处于不同的高度处。例如，在图像捕获装置122、124和126之间还可能存在侧向位移差，为处理单元110的立体分析给出额外视差信息。侧向位移差可记作d_x，如图2C和图2D所示。在一些实施方式中，前后位移(例如，范围位移)可能存在于图像捕获装置122、124和126之间。例如，图像捕获装置122可位于图像捕获装置124和/或图像捕获装置126后面的0.5米到2米以上。这种类型的位移可使图像捕获装置之一能够覆盖其它图像捕获装置的潜在盲点。

[0133] 图像捕获装置122可具有任何合适的分辨率功能(例如，与图像传感器关联的像素数目)，并且与图像捕获装置122关联的图像传感器的分辨率可大于、小于或等于与图像捕获装置124和126关联的图像传感器的分辨率。在一些实施方式中，与图像捕获装置122和/或图像捕获装置124和126关联的图像传感器可具有640x 480、1024x 768、1280x 960的分辨率或具有任何其它合适的分辨率。

IPR2025-00943
Tesla EX1002 Page 1267

[0134] 帧速率(例如,在移动到捕获与下一个图像帧关联的像素数据之前图像捕获装置获取一个图像帧的一组像素数据的速率)可以是可控的。与图像捕获装置122关联的帧速率可大于、小于或等于与图像捕获装置124和126关联的帧速率。与图像捕获装置122、124和126关联的帧速率可取决于可能影响帧速率的计时的各种因素。例如,图像捕获装置122、124和126中的一个或多个可包括在获取与图像捕获装置122、124和/或126中的图像传感器的一个或多个像素关联的图像数据之前或之后强加的可选像素延迟周期。一般,对应于每个像素的图像数据可根据装置的时钟速率获取(例如,每个时钟周期一个像素)。另外,在包括卷帘快门的实施方式中,图像捕获装置122、124和126中的一个或多个可包括在获取与图像捕获装置122、124和/或126中的图像传感器的一行像素关联的图像数据之前或之后强加的可选水平消隐周期。进一步,图像捕获装置122、124和/或126中的一个或多个可包括在获取与图像捕获装置122、124和126的图像帧关联的图像数据之前或之后强加的可选垂直消隐周期。

[0135] 这些计时控制可使与图像捕获装置122、124和126关联的帧速率同步,甚至当每个图像捕获装置的线扫描速率不同时亦如此。另外,如将在下面更详细地讨论的,在其它因素(例如,图像传感器分辨率、最大线扫描速率,等等)之中,这些可选计时控制可使图像捕获装置122的FOV与图像捕获装置124和126的一个或多个FOV重叠的区域的图像捕获同步,甚至当图像捕获装置122的视场不同于图像捕获装置124和126的FOV时亦如此。

[0136] 图像捕获装置122、124和126的帧速率计时可取决于关联的图像传感器的分辨率。例如,假设两个装置有类似的线扫描速率,如果一个装置包括分辨率为640x 480的图像传感器并且另一装置包括分辨率为1280x 960的图像传感器,则将需要更多时间来获取来自具有更高分辨率的传感器的图像数据的帧。

[0137] 可影响图像捕获装置122、124和126的图像数据获取的计时的另一因素是最大线扫描速率。例如,获取包括在图像捕获装置122、124和126中的图像传感器的一行图像数据将需要一些最小时间量。假设没有添加像素延迟周期,则获取一行图像数据的该最小时间量将与特定装置的最大线扫描速率相关。提供较高最大线扫描速率的装置潜在地提供比具有较小最大线扫描速率的装置更大的帧速率。在一些实施方式中,图像捕获装置124和126中的一个或多个的最大线扫描速率可大于与图像捕获装置122关联的最大线扫描速率。在一些实施方式中,图像捕获装置124和/或126的最大线扫描速率可以是图像捕获装置122的最大线扫描速率的1.25倍、1.5倍、1.75倍或2倍以上。

[0138] 在另一实施方式中,图像捕获装置122、124和126可具有相同的最大线扫描速率,但图像捕获装置122可在小于或等于其最大扫描速率的扫描速率下操作。所述系统可构造成使得图像捕获装置124和126中的一个或多个在等于图像捕获装置122的线扫描速率的线扫描速率下操作。在其它实例中,所述系统可构造成使得图像捕获装置124和/或图像捕获装置126的线扫描速率可以是图像捕获装置122的线扫描速率的1.25倍、1.5倍、1.75倍或2倍以上。

[0139] 在一些实施方式中,图像捕获装置122、124和126可以是非对称的。即,它们可包括具有不同视场(FOV)和焦距的相机。例如,图像捕获装置122、124和126的视场可包括相对于车辆200的环境的任何期望区域。在一些实施方式中,图像捕获装置122、124和126中的一个或多个可构造成获取来自车辆200前方、车辆200后方、车辆200侧方或它们的组合的环境的

IPR2025-00943
Tesla EX1002 Page 1268

图像数据。

[0140] 进一步，与每个图像捕获装置122、124和/或126关联的焦距可以是可选的(例如，通过包括适当的透镜等等)，使得每个装置均获取在相对于车辆200的期望距离范围处的对象的图像。例如，在一些实施方式中，图像捕获装置122、124和126可获取距车辆几米内的特写对象的图像。图像捕获装置122、124和126还可构造成获取更远离车辆的范围(例如，25m、50m、100m、150m以上)的对象的图像。进一步，图像捕获装置122、124和126的焦距可选择为使得一个图像捕获装置(例如，图像捕获装置122)可以获取相对接近车辆(例如，10m内或20m内)的对象的图像，而其它图像捕获装置(例如，图像捕获装置124和126)可以获取距车辆200更远的对象(例如，大于20m、50m、100m、150m，等等)的图像。

[0141] 根据一些实施方式，一个或多个图像捕获装置122、124和126的FOV可具有宽角度。例如，尤其对于可用于捕获车辆200附近的区域的图像的图像捕获装置122、124和126，可有利的是具有140度的FOV。例如，图像捕获装置122可用于捕获车辆200右方或左方的区域的图像，并且在这样的实施方式中，图像捕获装置122可期望的是具有宽FOV(例如，至少140度)。

[0142] 与图像捕获装置122、124和126中的每个关联的视场可取决于相应的焦距。例如，随着焦距的增加，对应的视场降低。

[0143] 图像捕获装置122、124和126可构造成具有任何合适的视场。在一个特定示例中，图像捕获装置122可具有46度的水平FOV，图像捕获装置124可具有23度的水平FOV，并且图像捕获装置126可具有介于23度到46度之间的水平FOV。在另一实例中，图像捕获装置122可具有52度的水平FOV，图像捕获装置124可具有26度的水平FOV，并且图像捕获装置126可具有26度到52度之间的水平FOV。在一些实施方式中，图像捕获装置122的FOV与图像捕获装置124和/或图像捕获装置126的FOV的比率可从1.5变化到2.0。在其它实施方式中，该比率可在1.25到2.25之间变化。

[0144] 系统100可构造成使得图像捕获装置122的视场与图像捕获装置124和/或图像捕获装置126的视场至少部分地或完全地重叠。在一些实施方式中，系统100可构造成使得图像捕获装置124和126的视场例如落入(例如，窄于)图像捕获装置122的视场内并且与图像捕获装置122的视场共享共同的中心。在其它实施方式中，图像捕获装置122、124和126可捕获相邻FOV，或者可使它们的FOV部分地重叠。在一些实施方式中，图像捕获装置122、124和126的视场可对准，使得较窄FOV图像捕获装置124和/或126的中心可位于较宽FOV装置122的视场的下半部中。

[0145] 图2F是与公开的实施方式相一致的示例性车辆控制系统的概略表示。如图2F中指示的，车辆200可包括节流系统220、制动系统230和转向系统240。系统100可经一个或多个数据链路(例如，用于发送数据的任何有线和/或无线的一个或多个链路)向节流系统220、制动系统230和转向系统240中的一个或多个提供输入(例如，控制信号)。例如，基于对由图像捕获装置122、124和/或126获取的图像的分析，系统100可向节流系统220、制动系统230和转向系统240中的一个或多个提供控制信号，以(例如，通过致使加速、转弯、车道移位，等等)导航车辆200。进一步，系统100可接收来自节流系统220、制动系统230和转向系统240中的一个或多个的输入，指示车辆200的操作条件(例如，速度，车辆200是否制动和/或转弯，等等)。下面，进一步细节结合图4至图7来提供。

IPR2025-00943
Tesla EX1002 Page 1269

[0146] 如图3A所示，车辆200还可包括用于与车辆200的驾驶员或乘客交互的用户接口170。例如，在车辆应用中，用户接口170可包括触摸屏320、旋钮330、按钮340和麦克风350。车辆200的驾驶员或乘客还可使用手柄(例如，位于车辆200的转向柱之上或附近，例如包括转向信号手柄)、按钮(例如，位于车辆200的方向盘上)等，以与系统100交互。在一些实施方式中，麦克风350可定位成与后视镜310相邻。类似地，在一些实施方式中，图像捕获装置122可位于后视镜310附近。在一些实施方式中，用户接口170还可包括一个或多个扬声器360(例如，车辆音频系统的扬声器)。例如，系统100可经由扬声器360提供各种通知(例如，提示)。

[0147] 图3B至图3D是与公开的实施方式相一致的示例性相机安装件370的图示，相机安装件370被构造成定位在后视镜(例如，后视镜310)后面并且抵着车辆挡风玻璃。如图3B所示，相机安装件370可包括图像捕获装置122、124和126。图像捕获装置124和126可定位在遮光板380后面，遮光板380可与车辆挡风玻璃齐平并且包括薄膜和/或抗反射材料的组合物。例如，遮光板380可定位成使得其抵着具有配合斜度的车辆挡风玻璃而对准。在一些实施方式中，例如，如图3D中描绘的，图像捕获装置122、124和126中的每个均可定位在遮光板380后面。公开的实施方式并不限于任何特定构造的图像捕获装置122、124和126、相机安装件370和遮光板380。图3C是图3B中示出的相机安装件370从正面视角的图示。

[0148] 如受益于本公开的本领域技术人员将了解的，对于前面的公开的实施方式可做出许多变型和/或修改。例如，并非所有部件对于系统100的操作都是至关重要的。进一步，任何部件都可位于系统100的任何适当的部分中，并且所述部件可重新布置成各种构造，同时提供公开的实施方式的功能性。因此，前面的构造是示例并且与上面讨论的构造无关，系统100可以提供宽范围的功能性，用以分析车辆200的周边环境并且响应于所述分析来导航车辆200。

[0149] 如下面进一步详细地讨论且与各种公开的实施方式相一致，系统100可提供与自主驱动和/或驾驶员辅助技术相关的各种特征。例如，系统100可分析图像数据、位置数据(例如，GPS地点信息)、地图数据、速度数据和/或来自包括在车辆200中的传感器的数据。系统100可收集来自例如图像获取单元120、位置传感器130及其它传感器的数据以供分析。进一步，系统100可分析所收集的数据以确定车辆200是否应该采取某些动作，然后在无人为干预的情况下自动采取确定的动作。例如，当车辆200在无人为干预的情况下导航时，系统100可(例如，通过向节流系统220、制动系统230和转向系统240中的一个或多个发送控制信号)自动控制车辆200的制动、加速和/或转向。进一步，系统100可分析所收集的数据并发出警告，和/或基于所收集的数据的分析来提示车辆驾乘人员。

[0150] 进一步，与公开的实施方式相一致，由系统100提供的功能性可致使车辆200采取不同的动作以在车道内和/或相对于其它车辆和/或对象来导航车辆200。例如，系统100可调整车辆200相对于车辆200行进的车道和/或相对于定位在车辆200附近的对象的定位，针对车辆200选择特定车道以在行进的同时使用，并且响应于侵入车辆(诸如试图进入车辆200行进的车道中的车辆)采取动作。另外，系统100可控制车辆200在不同场景(诸如当车辆200进行转弯时)下的速度。系统100可进一步致使车辆200模仿前车的动作或监测目标车辆，并且导航车辆200使之超越目标车辆。下面提供关于系统100所提供的各种实施方式的额外细节。

IPR2025-00943
Tesla EX1002 Page 1270

[0151] 朝前多成像系统

[0152] 如上面讨论的，系统100可提供使用多相机系统的驱动辅助功能性。多相机系统可使用面对车辆向前方向的一个或多个相机。在其它实施方式中，多相机系统可包括面对车辆侧面或车辆后面的一个或多个相机。在一个实施方式中，例如，系统100可使用双相机成像系统，其中第一相机和第二相机(例如，图像捕获装置122和124)可定位在车辆(例如，车辆200)的前面和/或侧面。第一相机的视场可大于、小于第二相机的视场，或者与第二相机的视场部分地重叠。另外，第一相机可连接到第一图像处理器，以执行第一相机所提供的图像的单眼图像分析，并且第二相机可连接到第二图像处理器，以执行第二相机所提供的图像的单眼图像分析。第一图像处理器和第二图像处理器的输出(例如，处理后的信息)可组合。在一些实施方式中，第二图像处理器可接收来自第一相机和第二相机两者的图像以执行立体分析。在另一实施方式中，系统100可使用三相机成像系统，其中每个相机均具有不同的视场。因此，这样的系统可基于从位于车辆前面和侧面两者的变化距离处的对象得到的信息做出决定。提到单眼图像分析可指这样的实例，即：基于从单个观察点(例如，从单个相机)捕获的图像来执行图像分析。立体图像分析可指这样的实例，即：基于用图像捕获参数的一个或多个变型捕获的两个或更多个图像来执行图像分析。例如，适于执行立体图像分析的被捕获图像可包括采取以下手段捕获的图像：根据两个或更多个不同位置，从不同的视场，使用不同的焦距，结合视差信息，等等。

[0153] 例如，在一个实施方式中，系统100可实施使用图像捕获装置122-126的三相机构造。在这样的构造中，图像捕获装置122可提供窄视场(例如，34度，或选自约20度到45度范围的其它值，等等)，图像捕获装置124可提供宽视场(例如，150度或选自约100度到约180度范围的其它值)，并且图像捕获装置126可提供中间视场(例如，46度或选自约35度到约60度范围的其它值)。在一些实施方式中，图像捕获装置126可作为主相机或主要相机。图像捕获装置122-126可定位在后视镜310后面并且基本上并排定位(例如，6cm间隔)。进一步，在一些实施方式中，如上面讨论的，图像捕获装置122-126中的一个或多个可安装在与车辆200的挡风玻璃齐平的遮光板380后面。这样的遮光可将来自车内的任何反射对图像捕获装置122-126造成的影响降到最低。

[0154] 在另一实施方式中，如上面结合图3B和图3C讨论的，宽视场相机(例如，上面示例中的图像捕获装置124)可安装得低于窄的主视场相机(例如，上面示例中的图像装置122和126)。该构造可提供来自宽视场相机的自由视线。为了减少反射，相机可安装得接近车辆200的挡风玻璃，并且可包括相机上的偏光器用以抑制反射光。

[0155] 三相机系统可提供某些性能特性。例如，一些实施方式可包括如下能力：基于另一相机的检测结果来验证一个相机针对对象的检测。在上面讨论的三相机构造中，处理单元110可包括例如三个处理装置(例如，如上面讨论的，三个EyeQ系列的处理器芯片)，每个处理装置均专用于处理由图像捕获装置122-126中的一个或多个捕获的图像。

[0156] 在三相机系统中，第一处理装置可接收来自主相机和窄视场相机两者的图像，并且执行窄FOV相机的视觉处理以例如检测其它车辆、行人、车道标线、交通标志、交通灯及其它道路对象。进一步，第一处理装置可计算主相机和窄相机的图像之间的像素差距并且创建车辆200的环境的三维重构。然后，第一处理装置可将三维重构与三维地图数据组合，或者与基于另一相机的信息计算的三维信息组合。

IPR2025-00943
Tesla EX1002 Page 1271

[0157] 第二处理装置可接收来自主相机的图像并且执行视觉处理以检测其它车辆、行人、车道标线、交通标志、交通灯及其它道路对象。另外，第二处理装置可计算相机位移，并且基于该位移计算连续图像之间的像素差距并创建场景的三维重构(例如，根据运动的结构)。第二处理装置可将来自基于运动的三维重构的结构发送到第一处理装置以与立体三维图像组合。

[0158] 第三处理装置可接收来自宽FOV相机的图像并且处理所述图像，以检测车辆、行人、车道标线、交通标志、交通灯及其它道路对象。第三处理装置可进一步执行额外处理指令以分析图像来识别在图像中移动的对象，诸如变道的车辆、行人，等等。

[0159] 在一些实施方式中，使基于图像的信息流被独立地捕获并处理可提供：提供系统冗余的机会。例如，这样的冗余可包括：使用第一图像捕获装置和从该装置处理的图像以验证和/或补充通过捕获并处理至少来自第二图像捕获装置的图像信息而获得的信息。

[0160] 在一些实施方式中，系统100可使用两个图像捕获装置(例如，图像捕获装置122和124)来为车辆200提供导航辅助，并且使用第三图像捕获装置(例如，图像捕获装置126)以提供冗余并验证从其它两个图像捕获装置接收的数据的分析。例如，在这样的构造中，图像捕获装置122和124可为系统100的立体分析提供图像以导航车辆200，而图像捕获装置126可为系统100的单眼分析提供图像以提供冗余并验证基于从图像捕获装置122和/或图像捕获装置124捕获的图像而获得的信息。即，图像捕获装置126(及对应的处理装置)可被认为提供冗余子系统，以提供对从图像捕获装置122和124得到的分析的检查(例如，提供自动紧急制动(AEB)系统)。

[0161] 本领域技术人员将辨识到，上面的相机构造、相机放置、相机数目、相机地点等等仅是示例。这些部件及相对于整体系统描述的其它部件可组装并用在各种不同构造中而不脱离公开的实施方式的范围。下面是关于多相机系统的使用的进一步细节，用以提供驾驶员辅助和/或自主车辆功能性。

[0162] 图4是与公开的实施方式相一致的存储器140和/或150的示例性功能框图，存储器140和/或150可存储/编程有用于执行一个或多个操作的指令。虽然以下是参照存储器140，但是本领域技术人员将辨识到，指令可存储在存储器140和/或150中。

[0163] 如图4所示，存储器140可存储单眼图像分析模块402、立体图像分析模块404、速度和加速度模块406以及导航响应模块408。公开的实施方式并不限于任何特定构造的存储器140。进一步，应用处理器180和/或图像处理器190可执行存储于包括在存储器140中的任何模块402-408之中的指令。本领域技术人员将理解，以下讨论中对处理单元110的引用可单独或共同地指应用处理器180和图像处理器190。因此，任何以下过程的步骤可由一个或多个处理装置执行。

[0164] 在一个实施方式中，单眼图像分析模块402可存储指令(诸如计算机视觉软件)，当由处理单元110执行时，该指令执行由图像捕获装置122、124和126之一获取的一组图像的单眼图像分析。在一些实施方式中，处理单元110可将来自一组图像的信息与额外的感觉信息(例如，来自雷达的信息)组合，以执行单眼图像分析。如下面结合图5A至图5D描述的，单眼图像分析模块402可包括用于检测一组图像内的一组特征(诸如车道标记、车辆、行人、道路标志、高速公路出口匝道、交通灯、危险对象及与车辆环境关联的任何其它特征)的指令。基于此分析，系统100(例如，经由处理单元110)可在车辆200中产生一个或多个导航响应，

IPR2025-00943
Tesla EX1002 Page 1272

诸如转弯、车道移位、加速度改变，等，这在下面结合导航响应模块408讨论。

[0165] 在一个实施方式中，立体图像分析模块404可存储指令(诸如计算机视觉软件)，当由处理单元110执行时，该指令执行通过选自任何图像捕获装置122、124和126的图像捕获装置的组合而获取的第一组和第二组图像的立体图像分析。在一些实施方式中，处理单元110可将来自第一组和第二组图像的信息与额外感觉信息(例如，来自雷达的信息)组合，以执行立体图像分析。例如，立体图像分析模块404可包括用于基于图像捕获装置124获取的第一组图像和图像捕获装置126获取的第二组图像来执行立体图像分析的指令。如下面结合图6描述的，立体图像分析模块404可包括用于检测第一组和第二组图像内的一组特征(诸如车道标记、车辆、行人、道路标志、高速公路出口匝道、交通灯、危险对象，等)的指令。基于此分析，处理单元110可在车辆200中产生一个或多个导航响应，诸如转弯、车道移位、加速度改变，等，这在下面结合导航响应模块408讨论。

[0166] 在一个实施方式中，速度和加速度模块406可存储构造成分析从车辆200的一个或多个计算机电装置接收的数据的软件，所述计算机电装置被构造成致使车辆200的速度和/或加速度改变。例如，处理单元110可执行与速度和加速度模块406关联的指令，以基于从执行单眼图像分析模块402和/或立体图像分析模块404得到的数据来计算车辆200的目标速度。例如，这样的数据可包括：目标位置、速度和/或加速度；车辆200相对于近旁车辆、行人或道路对象的位置和/或速度；车辆200相对于道路的车道标记的位置信息，等。另外，处理单元110可基于感觉输入(例如，来自雷达的信息)和来自车辆200的其它系统(诸如车辆200的节流系统220、制动系统230和/或转向系统240)的输入来计算车辆200的目标速度。基于计算出的目标速度，处理单元110可向车辆200的节流系统220、制动系统230和/或转向系统240发送电子信号，以通过例如物理踩下制动器或松开车辆200的加速器来触发速度和/或加速度的改变。

[0167] 在一个实施方式中，导航响应模块408可存储可由处理单元110执行的软件，以基于从单眼图像分析模块402和/或立体图像分析模块404的执行得到的数据来确定期望导航响应。这样的数据可包括与近旁车辆、行人和道路对象关联的位置和速度信息、车辆200的目标位置信息，等。另外，在一些实施方式中，导航响应可(部分地或完全地)基于地图数据、车辆200的预定位置和/或车辆200与从单眼图像分析模块402和/或立体图像分析模块404的执行而检测的一个或多个对象之间的相对速度或相对加速度。导航响应模块408还可基于感觉输入(例如，来自雷达的信息)和来自车辆200的其它系统(诸如车辆200的节流系统220、制动系统230和转向系统240)的输入来确定期望导航响应。基于所述期望导航响应，处理单元110可向车辆200的节流系统220、制动系统230和转向系统240发送电子信号，以触发期望导航响应，例如，通过使车辆200的方向盘转动以实现预定角度旋转。在一些实施方式中，处理单元110可使用导航响应模块408的输出(例如，期望导航响应)作为输入，以执行速度和加速度模块406来计算车辆200的速度改变。

[0168] 图5A是示出与公开的实施方式相一致的用于基于单眼图像分析来产生一个或多个导航响应的示例性过程500A的流程图。在步骤510，处理单元110可经由处理单元110与图像获取单元120之间的数据接口128接收多个图像。例如，包括在图像获取单元120中的相机(诸如具有视场202的图像捕获装置122)可捕获车辆200的前面(例如，或是车辆的侧面或后面)的区域的多个图像并且经数据连接(例如，数字、有线、USB、无线、蓝牙，等等)将图像发

IPR2025-00943
Tesla EX1002 Page 1273

送到处理单元110。在步骤520，处理单元110可执行单眼图像分析模块402以分析多个图像，这在下面结合图5B至图5D进一步详细地描述。通过执行所述分析，处理单元110可检测一组图像内的一组特征，诸如车道标记、车辆、行人、道路标志、高速公路出口匝道、交通灯，等。

[0169] 在步骤520，处理单元110还可执行单眼图像分析模块402以检测各种道路危险，例如，卡车轮胎的一部分、跌落的道路标志、松动的货物、小动物，等。道路危险可能存在结构、形状、大小和颜色上的变化，可能使检测这样的危险更具挑战性。在一些实施方式中，处理单元110可执行单眼图像分析模块402，以对多个图像进行多帧分析来检测道路危险。例如，处理单元110可估算连贯图像帧之间的相机运动，并且计算帧之间的像素差距以构建道路的三维地图。然后，处理单元110可使用三维地图来检测路面以及存在于路面上方的危险。

[0170] 在步骤530，处理单元110可执行导航响应模块408，以基于步骤520执行的分析以及如上面结合图4描述的技术来产生车辆200的一个或多个导航响应。例如，导航响应可包括转弯、车道移位、加速度改变，等。在一些实施方式中，处理单元110可使用从执行速度和加速度模块406得到的数据，以产生一个或多个导航响应。另外，多个导航响应可同时、顺次或按其任何组合而发生。例如，处理单元110可致使车辆200变换一个车道，然后通过例如向车辆200的转向系统240和节流系统220依次发送控制信号来加速。替代地，处理单元110可致使车辆200制动，同时通过例如向车辆200的制动系统230和转向系统240同时发送控制信号来使变道。

[0171] 图5B是示出与公开的实施方式相一致的用于检测一组图像中的一个或多个车辆和/或行人的示例性过程500B的流程图。处理单元110可执行单眼图像分析模块402以实施过程500B。在步骤540，处理单元110可确定表示可能的车辆和/或行人的一组候选对象。例如，处理单元110可扫描一个或多个图像，将所述图像与一个或多个预定图案比较，并且在每个图像内识别可包含关注对象(例如，车辆、行人，或其一部分)的可能地点。预定图案可采取实现高比率“错误击中”和低比率“未击中”这样的方式来设计。例如，处理单元110可使用与预定图案低阈值的相似度来识别候选对象作为可能的车辆或行人。这样做可允许处理单元110减少未击中(例如，未识别)表示车辆或行人的候选对象的概率。

[0172] 在步骤542，处理单元110可过滤这一组候选对象以基于分类标准来排除某些候选(例如，无关或不太相关的对象)。这样的标准可从与存储在数据库(例如，存储在存储器140中的数据库)中的对象类型关联的各种属性得到。属性可包括对象形状、尺寸、质地、(例如，相对于车辆200的)位置，等。由此，处理单元110可使用一组或多组标准以拒绝来自一组候选对象的错误候选。

[0173] 在步骤544，处理单元110可分析图像的多个帧，以确定一组候选对象中的对象是否表示车辆和/或行人。例如，处理单元110可跨越连贯的帧来跟踪被检测候选对象并且积累与被检测对象关联的逐帧数据(例如，大小、相对于车辆200的位置，等等)。另外，处理单元110可估算被检测对象的参数并且将对象的逐帧位置数据与预测位置比较。

[0174] 在步骤546，处理单元110可构建被检测对象的一组测量结果。例如，这样的测量结果可包括与被检测对象关联的(相对于车辆200的)位置、速度和加速度值。在一些实施方式中，处理单元110可基于以下技术来构建测量结果：基于估算技术使用一系列基于时间的观察(诸如卡尔曼滤波器或线性二次估算(LQE))；和/或基于不同对象类型(例如，汽车、卡车、行人、自行车、道路标志，等等)的可用建模数据。卡尔曼滤波器可基于对象尺度的测量结

IPR2025-00943
Tesla EX1002 Page 1274

果，其中尺度测量结果与将要碰撞的时间(例如，车辆200到达对象的时间量)成比例。由此，通过执行步骤540-546，处理单元110可识别出现在一组被捕获图像内的车辆和行人并且得到与车辆和行人关联的信息(例如，位置、速度、大小)。基于识别和得到的信息，处理单元110可产生车辆200的一个或多个导航响应，这在上面结合图5A描述。

[0175] 在步骤548，处理单元110可执行一个或多个图像的光流分析，以减少检测“错误击中”以及未击中表示车辆或行人的候选对象的概率。光流分析可指的是，例如，分析相对于与其它车辆和行人关联的一个或多个图像中的车辆200的运动图案，并且区别于路面运动。处理单元110可通过观察跨越多个图像帧(在不同的时间被捕获)的对象的不同位置来计算候选对象的运动。处理单元110可使用位置和时间值作为用于计算候选对象运动的数学模型的输入。由此，光流分析可提供检测在车辆200近旁的车辆和行人的另一种方法。处理单元110可与步骤540-546组合来执行光流分析，为检测车辆和行人提供冗余并且增加系统100的可靠性。

[0176] 图5C是示出与公开的实施方式相一致的用于检测一组图像中的路标和/或车道几何信息的示例性过程500C的流程图。处理单元110可执行单眼图像分析模块402以实施过程500C。在步骤550，处理单元110可通过扫描一个或多个图像来检测一组对象。为了检测一段车道标记、车道几何信息及其它相关路标，处理单元110可过滤一组对象以排除确定无关的这些对象(例如，轻微的坑洼、小石头，等等)。在步骤552，处理单元110可将步骤550中检测的属于相同路标或车道标线的区段分组到一起。基于该分组，处理单元110可建立表示被检测区段的模型，诸如数学模型。

[0177] 在步骤554，处理单元110可构建与被检测区段关联的一组测量结果。在一些实施方式中，处理单元110可创建被检测区段从图像平面向现实世界平面的投影。该投影可使用第三多项式来表征，该第三多项式具有对应于物理属性(诸如位置、斜度、曲率和被检测道路的曲率衍生)的系数。在生成投影时，处理单元110可考虑路面改变，以及与车辆200关联的俯仰和横摆速率。另外，处理单元110可通过分析存在于路面上的位置和运动线索来建模道路标高。进一步，处理单元110可通过跟踪一个或多个图像中的一组特征点来估算与车辆200关联的俯仰和横摆速率。

[0178] 在步骤556，处理单元110可通过例如跟踪跨越连贯图像帧的被检测区段以及积累与被检测区段关联的逐帧数据来执行多帧分析。在处理单元110执行多帧分析时，在步骤554构建的一组测量结果可变得更可靠并且与更高的置信水平关联。由此，通过执行步骤550-556，处理单元110可识别出现在一组被捕获图像内的路标并且得到车道几何信息。基于该识别和得到的信息，处理单元110可产生车辆200的一个或多个导航响应，这在上面结合图5A描述。

[0179] 在步骤558，处理单元110可考虑额外信息源来进一步在其周边环境的背景下建立车辆200的安全模型。处理单元110可使用安全模型来限定以下背景，即：系统100可采取安全的方式来执行车辆200的自主控制。为了建立安全模型，在一些实施方式中，处理单元110可考虑其它车辆的位置和运动、被检测道路的路边和障碍和/或从地图数据(诸如地图数据库160的数据)提取的总体道路形状描述。通过考虑额外信息源，处理单元110可提供冗余来检测路标和车道几何形状并且增加系统100的可靠性。

[0180] 图5D是示出与公开的实施方式相一致的用于检测一组图像中的交通灯的示例性

IPR2025-00943
Tesla EX1002 Page 1275

过程500D的流程图。处理单元110可执行单眼图像分析模块402以实施过程500D。在步骤560，处理单元110可扫描一组图像并且识别出现在可能包含交通灯的图像中的地点处的对象。例如，处理单元110可过滤被识别对象以构建一组候选对象，排除不可能对应于交通灯的这些对象。过滤可基于与交通灯关联的各种属性(诸如形状、尺寸、质地、(例如，相对于车辆200的)位置，等)来完成。这样的属性可基于交通灯和交通控制信号的多个示例并且存储在数据库中。在一些实施方式中，处理单元110可在反映可能交通灯的一组候选对象上执行多帧分析。例如，处理单元110可跟踪跨越连贯图像帧的候选对象，估算候选对象的现实世界位置，并且过滤出移动的那些对象(不可能是交通灯)。在一些实施方式中，处理单元110可执行候选对象的颜色分析，并且识别出现在可能交通灯内侧的被检测颜色的相对位置。

[0181] 在步骤562，处理单元110可分析枢纽的几何形状。该分析可基于以下的任何组合：(i)在车辆200的任一侧的检测的车道数目；(ii)在道路上检测的标记(诸如箭头标线)；以及(iii)从地图数据(诸如地图数据库160的数据)提取的枢纽的描述。处理单元110可使用从执行单眼分析模块402得到的信息来进行分析。另外，处理单元110可确定在步骤560检测的交通灯与出现在车辆200附近的车道之间的对应性。

[0182] 在车辆200接近枢纽时，在步骤564，处理单元110可更新与被分析的枢纽几何形状和被检测交通灯关联的置信水平。例如，与实际出现在枢纽处的交通灯数目比较，估算的出现在枢纽处的交通灯数目可能影响置信水平。由此，基于该置信水平，处理单元110可授权车辆200的驾驶员进行控制，以便改进安全条件。通过执行步骤560-564，处理单元110可识别出现在一组被捕获图像内的交通灯并且分析枢纽几何形状信息。基于该识别以及分析，处理单元110可产生车辆200的一个或多个导航响应，这在上面结合图5A描述。

[0183] 图5E是示出与公开的实施方式相一致的用于基于车辆路径来在车辆200中产生一个或多个导航响应的示例性过程500E的流程图。在步骤570，处理单元110可构建与车辆200关联的初始车辆路径。车辆路径可使用以坐标(x,z)表达的一组点来表示，并且一组点的两点之间的距离d_i可落入1米到5米的范围内。在一个实施方式中，处理单元110可使用诸如左右道路多项式的两个多项式来构建初始车辆路径。处理单元110可计算两个多项式之间的几何中点，并且如果有偏移的话，使包括在所得的车辆路径中的每个点偏移预定偏移量(例如，智能车道偏移)，(零偏移可对应于行进在车道的中间)。偏移的方向可垂直于车辆路径中的任何两点之间的区段。在另一实施方式中，处理单元110可使用一个多项式和估算的车道宽度使车辆路径的每个点偏移达：估算的车道宽度的一半加上预定偏移量(例如，智能车道偏移)。

[0184] 在步骤572，处理单元110可更新在步骤570构建的车辆路径。处理单元110可使用更高分辨率来重新构建在步骤570构建的车辆路径，使得表示车辆路径的一组点的两点之间的距离d_k小于上面描述的距离d_i。例如，距离d_k可落入0.1米到0.3米的范围内。处理单元110可使用抛物面样条算法来重新构建车辆路径，这可产生对应于车辆路径的总长度的累计距离向量S(即，基于表示车辆路径的一组点)。

[0185] 在步骤574，处理单元110可基于在步骤572构建的更新车辆路径来确定前瞻点(以坐标(x_l,z_l)表达)。处理单元110可从累计距离向量S提取前瞻点，并且前瞻点可与前瞻距离和前瞻时间关联。下限范围可从10米到20米的前瞻距离可计算为车辆200的速度和前瞻时间的乘积。例如，随着车辆200的速度降低，前瞻距离也可减小(例如，直到前瞻距离达到下

IPR2025-00943
Tesla EX1002 Page 1276

限)。范围可从0.5秒到1.5秒的前瞻时间可与在车辆200中产生导航响应相关联的一个或多个控制环的增益(诸如跟踪控制环的航向误差)成反比。例如,跟踪控制环的航向误差的增益可取决于偏航率环的带宽、转向致动器环、汽车侧向动态,等。由此,跟踪控制环的航向误差的增益越大,前瞻时间越短。

[0186] 在步骤576,处理单元110可基于在步骤574确定的前瞻点来确定航向误差和偏航率命令。处理单元110可通过计算前瞻点的反正切(例如,$\arctan(x_1/z_1)$)来确定航向误差。处理单元110可将偏航率命令确定为航向误差和高水平控制增益的乘积。如果前瞻距离未处于下限,则高水平控制增益可等于:(2/前瞻时间)。以其它方式,高水平控制增益可等于:(车辆200的2倍速度/前瞻距离)。

[0187] 图5F是示出与公开的实施方式相一致的用于确定前车是否变道的示例性过程500F的流程图。在步骤580,处理单元110可确定与前车(例如,行进在车辆200之前的车辆)关联的导航信息。例如,处理单元110可使用上面结合图5A和图5B描述的技术来确定前车的位置、速度(例如,方向和速度)和/或加速度。处理单元110还可使用上面结合图5E描述的技术来确定一个或多个道路多项式、(与车辆200关联的)前瞻点和/或缓慢移动踪迹(例如,描述前车采取的路径的一组点)。

[0188] 在步骤582,处理单元110可分析在步骤580确定的导航信息。在一个实施方式中,处理单元110可计算缓慢移动踪迹与道路多项式(例如,沿着踪迹)之间的距离。如果沿着踪迹的该距离的方差超过预定阈值(例如,笔直道路上为0.1米到0.2米,适度弯曲的道路上为0.3米到0.4米,并且急剧弯曲的道路上为0.5米到0.6米),处理单元110可确定:前车可能变道。在检测有多个车辆行进在车辆200之前的情况下,处理单元110可比较与每个车辆关联的缓慢移动踪迹。基于该比较,处理单元110可确定:缓慢移动踪迹与其它车辆的缓慢移动踪迹不匹配的车辆可能变道。处理单元110可另外将(与前车关联的)缓慢移动踪迹的曲率与前车行进的道路区段的预期曲率比较。预期曲率可从地图数据(例如,地图数据库160的数据)、道路多项式、其它车辆的缓慢移动踪迹、关于道路的现有知识等之中提取。如果缓慢移动踪迹与道路区段的预期曲率的曲率差超过预定阈值,则处理单元110可确定:前车可能变道。

[0189] 在另一实施方式中,处理单元110可在具体时间段(例如,0.5秒到1.5秒)内将前车的瞬时位置与(与车辆200关联的)前瞻点比较。如果前车瞬时位置与前瞻点之间的距离在具体时间段期间变化,并且变化的累计总和超过预定阈值(例如,笔直道路上为0.3米到0.4米,适度弯曲的道路上为0.7米到0.8米,并且急剧弯曲的道路上为1.3米到1.7米),则处理单元110可确定:前车可能变道。在另一实施方式中,处理单元110可通过将沿着踪迹行进的侧向距离与缓慢移动踪迹的预期曲率比较来分析缓慢移动踪迹的几何形状。预期曲率半径可根据下面的计算来确定:$(\delta_z^2+\delta_x^2)/2/(\delta_x)$,其中$\delta_x$表示行进的侧向距离,并且$\delta_z$表示行进的纵向距离。如果行进的侧向距离与预期曲率之间的差超过预定阈值(例如,500米到700米),则处理单元110可确定:前车可能变道。在另一实施方式中,处理单元110可分析前车的位置。如果前车的位置掩盖了道路多项式(例如,前车覆盖在道路多项式的顶部),则处理单元110可确定:前车可能变道。在前车的位置使得另一车辆被检测位于前车之前并且两个车辆的缓慢移动踪迹不平行的情况下,处理单元110可确定:(更接近的)前车可能变道。

[0190] 在步骤584,处理单元110可基于在步骤582执行的分析来确定前车200是否变道。

IPR2025-00943
Tesla EX1002 Page 1277

例如，处理单元110可基于在步骤582执行的各个分析的加权平均来做出确定。在这样的方案之下，例如，处理单元110基于特定类型的分析做出的前车可能变道的决定可被分配值“1”（而“0”表示确定前车不可能变道）。在步骤582执行的不同分析可被分配不同的权重，并且公开的实施方式并不限于任何特定组合的分析和权重。

[0191] 图6是示出与公开的实施方式相一致的用于基于立体图像分析来产生一个或多个导航响应的示例性过程600的流程图。在步骤610，处理单元110可经由数据接口128接收第一和第二多个图像。例如，包括在图像获取单元120中的相机（诸如具有视场202和204的图像捕获装置122和124）可捕获车辆200前方的区域的第一和第二多个图像并且经数字连接（例如，USB、无线、蓝牙，等等）将第一和第二多个图像发送到处理单元110。在一些实施方式中，处理单元110可经由两个或更多个数据接口接收第一和第二多个图像。公开的实施方式并不限于任何特定的数据接口构造或协议。

[0192] 在步骤620，处理单元110可执行立体图像分析模块404来执行第一和第二多个图像的立体图像分析，以创建车辆前方道路的三维地图并且检测图像内的特征，诸如车道标记、车辆、行人、道路标志、高速公路出口匝道、交通灯、道路危险，等。立体图像分析可采取类似于上面结合图5A至图5D描述的步骤的方式执行。例如，处理单元110可执行立体图像分析模块404以检测第一和第二多个图像内的候选对象（例如，车辆、行人、路标、交通灯、道路危险，等等），基于各种标准来过滤出候选对象的子集，并且执行多帧分析、构建测量结果并确定其余候选对象的置信水平。在执行上面的步骤时，处理单元110可考虑来自第一和第二多个图像两者的信息，而非单独来自一组图像的信息。例如，处理单元110可针对出现在第一和第二多个图像两者中的候选对象分析像素水平数据（或来自两个被捕获图像流之中的其它数据子集）的差别。作为另一示例，处理单元110可通过观察对象出现在多个图像中的一个之中但不出现在另一个中或者相对于出现在两个图像流中的对象可能存在的其它差别来估算候选对象（例如，相对于车辆200）的位置和/或速度。例如，相对于车辆200的位置、速度和/或加速度可基于与出现在图像流的一者或两者中的对象关联的特征的轨迹、位置、移动特性等等来确定。

[0193] 在步骤630，处理单元110可基于在步骤620执行的分析和上面结合图4描述的技术来执行导航响应模块408，以在车辆200中产生一个或多个导航响应。例如，导航响应可包括转弯、车道移位、加速度改变、速度改变、制动，等。在一些实施方式中，处理单元110可使用从执行速度和加速度模块406中得到的数据，以产生一个或多个导航响应。另外，多个导航响应可同时、顺次或按其任何组合而发生。

[0194] 图7是示出与公开的实施方式相一致的用于基于对三组图像的分析来产生一个或多个导航响应的示例性过程700的流程图。在步骤710，处理单元110可经由数据接口128接收第一、第二和第三多个图像。例如，包括在图像获取单元120中的相机（诸如具有视场202、204和206的图像捕获装置122、124和126）可捕获车辆200前方和/或侧面的区域的第一、第二和第三多个图像并且经数字连接（例如，USB、无线、蓝牙，等等）将第一、第二和第三多个图像发送到处理单元110。在一些实施方式中，处理单元110可经由三个以上数据接口接收第一、第二和第三多个图像。例如，图像捕获装置122、124、126中的每个均可具有关联的数据接口用于与处理单元110通信数据。公开的实施方式并不限于任何特定的数据接口构造或协议。

IPR2025-00943
Tesla EX1002 Page 1278

[0195] 在步骤720，处理单元110可分析第一、第二和第三多个图像以检测图像内的特征，诸如车道标记、车辆、行人、道路标志、高速公路出口匝道、交通灯、道路危险，等。该分析可采取类似于上面结合图5A至图5D和图6描述的步骤的方式执行。例如，处理单元110可对第一、第二和第三多个图像中的每个执行单眼图像分析（例如，经由执行单眼图像分析模块402并且基于上面结合图5A至图5D描述的步骤）。替代地，处理单元110可对第一和第二多个图像、第二和第三多个图像和/或第一和第三多个图像执行立体图像分析（例如，经由执行立体图像分析模块404并且基于上面结合图6描述的步骤）。对应于第一、第二和/或第三多个图像的分析的处理后的信息可加以组合。在一些实施方式中，处理单元110可执行单眼图像分析和立体图像分析的组合。例如，处理单元110可对第一多个图像执行单眼图像分析（例如，经由执行单眼图像分析模块402）以及对第二和第三多个图像执行立体图像分析（例如，经由执行立体图像分析模块404）。图像捕获装置122、124和126的构造（包括其相应地点和视场202、204和206）可影响对第一、第二和第三多个图像进行的分析的类型。公开的实施方式并不限于特定构造的图像捕获装置122、124和126，或者对第一、第二和第三多个图像进行的分析的类型。

[0196] 在一些实施方式中，处理单元110可基于在步骤710和720获取并分析的图像来执行系统100的测试。这样的测试可向某些构造的图像捕获装置122、124和126提供系统100的整体性能的指示器。例如，处理单元110可确定“错误击中”（例如，系统100不正确地确定存在车辆或行人的情况）和“未击中”的比例。

[0197] 在步骤730，处理单元110可基于从第一、第二和第三多个图像中的两者得到的信息来在车辆200中产生一个或多个导航响应。选择第一、第二和第三多个图像中的两者可取决于各种因素，例如，多个图像中的每个中检测的对象的数目、类型和大小。处理单元110还可基于以下做出选择：图像质量和分辨率、图像中反射的有效视场、被捕获帧的数目、关注的一个或多个对象实际出现在帧中的程度（例如，对象出现的帧的百分比、出现在每个这样的帧中的对象的比例，等等），等。

[0198] 在一些实施方式中，处理单元110可通过确定从一个图像源得到的信息与从其它图像源得到的信息的一致程度来选择从第一、第二和第三多个图像中的两者得到的信息。例如，处理单元110可组合从图像捕获装置122、124和126中的每个得到的被处理信息（不论通过单眼分析、立体分析或两者的任何组合）并且确定跨越从图像捕获装置122、124和126中的每个捕获的图像呈现一致的视觉指示器（例如，车道标记、被检测车辆及其地点和/或路径、被检测交通灯，等等）。处理单元110还可排除跨越被捕获图像不一致的信息（例如，变道的车辆，指示太接近车辆200的车辆的车道模型，等等）。由此，处理单元110可基于一致和不一致信息的确定来选择从第一、第二和第三多个图像中的两者得到的信息。

[0199] 例如，导航响应可包括转弯、车道移位、加速度改变，等。处理单元110可基于在步骤720执行的分析以及上面结合图4描述的技术来产生一个或多个导航响应。处理单元110还可使用从执行速度和加速度模块406得到的数据，以产生一个或多个导航响应。在一些实施方式中，处理单元110可基于车辆200与任何第一、第二和第三多个图像内的被检测对象之间的相对位置、相对速度和/或相对加速度来产生一个或多个导航响应。多个导航响应可同时、顺次或按其任何组合而发生。

[0200] 智能车道偏移

IPR2025-00943
Tesla EX1002 Page 1279

[0201] 系统100可提供驾驶员辅助功能性，以监测并且调整车辆200与车辆200的行进边界（例如，车道边界）之间的自由空间。车道可指车辆的指定或预期行进路径并且可具有标记边界（例如，道路上的线）或未标记边界（例如，道路路边、道路障碍、护栏、停放的车辆，等等）。例如，默认地，系统100可最大化或增加车辆200与当前车道边界之间的自由空间。进一步，在某些条件下，系统100可使车辆200朝向当前行进车道的一侧偏移而不离开车道。例如，在车辆200围绕弯道行进的情况下，系统100可降低朝向弯道内侧的偏移（例如，调整车辆200的位置使之更接近弯道的内侧）。作为另一示例，在障碍（诸如停放的车、行人或骑自行车的人）存在于车道一侧的情况下，系统100可在车辆200的存在障碍的那一侧增加偏移（例如，调整车辆200的位置使得更远离障碍）。

[0202] 图8A图示了行进在道路800上的车辆200，其中可使用用于识别车道约束并且在车道约束内操作车辆200的公开的系统和方法。如图所示，车辆200可具有第一车辆侧802，第一车辆侧802距第一车道约束810可以是第一距离805。类似地，车辆200可具有与第一车辆侧802相反的第二车辆侧812，并且第二车辆侧812距第二车道约束820可以是第二距离815。以这种方式，第一车道约束810和第二车道约束820可限定可供车辆200行进的车道。

[0203] 处理单元110可构造成基于处理单元110可经由数据接口128接收的由图像捕获装置122-126获取的多个图像来确定第一车道约束810和第二车道约束820。根据一些实施方式，第一车道约束810和/或第二车道约束820可由可见的车道边界（诸如标记在路面上的虚线或实线）识别。另外地或替代地，第一车道约束810和/或第二车道约束820可包括路面的路边或障碍。另外地或替代地，第一车道约束810和/或第二车道约束820可包括标记物（例如，博茨点）。根据一些实施方式，处理单元110可通过识别路面宽度的中点来确定第一车道约束810和/或第二车道约束820。根据一些实施方式，如果处理单元110仅识别第一车道约束810，则处理单元110可诸如基于估算的车道宽度或道路宽度来估算第二车道约束820。当例如指定道路车道的线未被喷涂或以其它方式贴标签时，处理单元110可采取这种方式识别车道约束。

[0204] 检测第一车道约束810和/或第二车道约束820可包括：处理单元110确定其在相机坐标系统中的三维模型。例如，第一车道约束810和/或第二车道约束820的三维模型可由第三多项式描述。除了三维建模第一车道约束810和/或第二车道约束820，处理单元110可执行主机运动参数（诸如车辆200的速度、偏航和俯仰速率以及加速度）的多帧估算。可选地，处理单元110可检测静态和移动车辆及其位置、航向、速度和加速度（都相对于车辆200）。处理单元110可进一步确定道路标高模型以将从多个图像获取的所有信息转换为三维空间。

[0205] 一般，作为默认条件，车辆200可在第一车道约束和第二车道约束810和/或820内相对居中地行进。然而，在一些情形下，环境因素可使之非期望或不可行，诸如当对象或结构存在于道路的一侧的时候。由此，可存在以下情形，即车辆200可期望或有利地更接近一个车道约束或另一个。处理单元110还可构造成确定被称为车道偏移条件的这样的情形是否基于多个图像存在。

[0206] 为了确定车道偏移条件是否存在，处理单元110可构造成确定一个或多个对象825存在于车辆200附近。例如，对象825可包括在道路800上行进的另一车辆，诸如汽车、卡车或摩托车。对于每个对象825，处理单元110可确定偏移轮廓。偏移轮廓可包括确定在其当前预测位置中对象825是否将位于车辆200的预定范围内。如果处理单元110确定对象825位于或

IPR2025-00943
Tesla EX1002 Page 1280

将位于车辆200的预定范围内，则处理单元110可确定：为使车辆200绕过对象825，在第一车道约束810和第二车道约束820内是否有足够的空间。如果没有足够的空间，则处理单元110可执行变道。如果有足够的空间，则处理单元110可确定在车辆200与对象825之间的距离接近之前是否有足够的时间绕过对象825。如果有足够的时间，则处理单元110可启动或激活偏移机动以绕过对象825。偏移机动可由处理单元110执行，使得车辆200的移动对驾驶员而言平稳而舒适。例如，偏移摆率可大约为0.15米/秒到0.75米/秒。偏移机动可执行为使得偏移机动的最大振幅将发生在车辆200与对象825之间的间距缩短的时候。如果没有足够的时间，则处理单元110可启动制动(例如，通过向制动系统230发送电子信号)，然后启动偏移机动。

[0207] 在为每个对象825设定好每个偏移轮廓均之后，处理单元可组合偏移轮廓，使得与车辆200绕过的每个对象825保持最大侧向距离。即，偏移机动可被修改和/或执行为使得车辆200与车辆200绕过的每个对象825之间的侧向距离被最大化。

[0208] 另外地或替代地，对象825可以是停放或静止的车辆、墙或人(诸如行人)或自行车交通。处理单元110可确定：基于对象825的某些特性存在车道偏移条件。例如，处理单元110可确定：如果对象825的高度超过预定阈值，诸如距路面10cm，则对象825构成车道偏移条件。以这种方式，处理单元110可构造成过滤出将不会出现车道偏移条件的对象，诸如小件碎片。

[0209] 另外地或替代地，处理单元110可构造成基于目标对象830在车辆第一侧802的位置来确定车道偏移条件是否存在于第一车辆侧802。目标对象830可包括构成车道偏移条件的一个或多个对象。另外地或替代地，目标对象830可不构成车道偏移条件。例如，目标对象830可以是街道标志或足够远离行进车道或足以不频繁发生的其它对象(不构成车道偏移条件)。处理单元110可构造成如果这样的目标对象830被检测则得出结论车道偏移条件不存在。

[0210] 处理单元110可构造成致使车辆200行进，使得第一距离805和第二距离810基本不同，以避免或保持更远离对象825。例如，如果车道偏移条件存在于第一车辆侧802，则处理单元110可构造成致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805大于第二距离815。这例如在图8A中图示。替代地，如果车道偏移条件存在于第二车辆侧812上，则处理单元110可创建或执行这样的偏移轮廓，该偏移轮廓被构造成致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805小于第二距离815。这例如在图8B中图示。

[0211] 作为另一示例，停放的车辆可位于车辆200的一侧，并且行人可位于车辆200的另一侧。在这样的情形下，系统100可确定：期望减小车辆200与停放的车辆之间的距离，以便在车辆200与行人之间提供更多空间。因此，处理单元110可致使车辆200行进，使得停放的车辆之间的距离减小。

[0212] 第一车辆侧802和第二车辆侧812的识别需要不是静态的，使得在不同的时间点，处理单元110可将第一车辆侧802识别为车辆200的驾驶员侧或乘客侧。例如，取决于一个或多个对象825存在于车辆200的哪一侧，第一车辆侧802的识别可改变，使得如果基于对象825存在一个车道偏移条件，则对象825存在于第一车辆侧802上。在这样的情形下，处理单元110可创建或执行这样的偏移轮廓，该偏移轮廓被构造成致使车辆200在第一车道约束

IPR2025-00943
Tesla EX1002 Page 1281

810和第二车道约束820内行进，使得大于第二距离815的第一距离805被朝向对象825提供。

[0213] 根据一些实施方式，当车道偏移条件存在于第一车辆侧802时，处理单元110可创建这样的偏移轮廓，该偏移轮廓被构造成致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805比第二距离815大至少1.25倍。另外地或替代地，当车道偏移条件存在于第一车辆侧802上时，处理单元110可创建这样的偏移轮廓，该偏移轮廓被构造成致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805比第二距离815大至少2倍。

[0214] 在一些实施方式中，如图8C所示，处理单元110可确定车道偏移条件存在于第一车辆侧802和第二车辆侧812两者上。然后，处理单元110可针对该偏移条件创建这样的偏移轮廓，该偏移轮廓被构造成致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805基本上等于第二距离815。

[0215] 处理单元110可构造成使得用于移动对象825以及静止对象825的偏移轮廓优先。例如，处理单元110可优先偏移轮廓，以保持车辆200与移动对象825和/或静态对象825之间的预定距离。处理单元200可基于某些车道特性(诸如弯道)优先这些偏移条件(优于这些偏移条件)。最后优先车道偏移条件可包括模仿前车。

[0216] 处理单元110可将偏移轮廓或偏移条件化为期望的前瞻点，该期望的前瞻点可位于由第一车道约束810和第二车道约束820限定的车道的中间处。处理单元110可基于前瞻点和偏移轮廓来计算并实施偏航率命令。偏航率命令可用于基于偏航率命令和车辆200的通过使用陀螺仪测量的偏航率来计算扭矩命令。处理单元110可向车辆200的转向系统(例如，转向系统240)发送扭矩命令。

[0217] 根据一些实施方式，诸如经由用户接口170，来自车辆200的驾驶员的输入可撤销处理单元110。例如，用户输入可撤销处理单元110的控制，致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805大于第二距离815。

[0218] 另外地或替代地，处理单元110可构造成基于由图像捕获装置122捕获的多个图像来确定车道偏移条件是否基本上不存在。这可包括以下情形，其中低于某些大小阈值的一个或多个对象825被检测，和/或当无对象825被检测时，被检测的一个或多个对象825位于距车道约束的预定距离的外侧。在这样的情形下，处理单元110可构造成致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805等于或基本上等于第二距离815。

[0219] 如图8D中图示的，车道偏移条件可包括弯路。处理单元110可构造成致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805相对于车辆200朝向弯路的外侧，并且第二距离815相对于车辆200朝向弯路的内侧。例如，第一距离805可小于第二距离815。

[0220] 图9是与公开的实施方式相一致的存储器140和/或150的示例性框图，存储器140和/或150可存储用于执行一个或多个操作的指令。如图9中图示的，存储器140可存储用于执行本文中描述的车道偏移条件检测和响应的一个或多个模块。例如，存储器140可存储车道约束模块910和车道偏移模块920。应用处理器180和/或图像处理器190可执行存储在任何模块910和920(包括在存储器140中)中的指令。本领域技术人员将理解的是，以下讨论对处理单元110的引用可单独或共同指应用处理器180和图像处理器190。因此，任何以下过程

IPR2025-00943
Tesla EX1002 Page 1282

的步骤可由一个或多个处理装置执行。

[0221] 车道约束模块910可存储指令，当由处理单元110执行时，该指令可检测并且限定第一车道约束810和第二车道约束820。例如，车道偏移模块910可处理从至少一个图像捕获装置122-124接收的多个图像，以检测第一车道约束810和第二车道约束820。如上面讨论的，这可包括识别喷涂的车道线和/或测量路面的中点。

[0222] 车道偏移模块920可存储指令，当由处理单元110执行时，该指令可检测车道偏移条件的存在和/或识别车道偏移条件是否存在于第一车辆侧802和/或第二车辆侧812上。例如，车道偏移模块920可处理多个图像，以检测对象825存在于第一车辆侧802和/或第二车辆侧812上。车道偏移模块920可处理多个图像，以检测车辆200行进的道路中的弯道。另外地或替代地，车道偏移模块920可接收来自指示存在对象825和/或弯曲车道的另一模块或其它系统的信息。车道偏移模块920可执行控制，以改变车辆200相对于第一车道约束810和/或第二车道约束820的位置。例如，如果车道偏移模块920确定车道偏移条件仅在第一车辆侧802上，则车道偏移模块920可使车辆200移动得更接近第二车道约束820。另一方面，如果车道偏移模块920确定车道偏移条件仅在第一车辆侧812上，则车道偏移模块920可使车辆200移动得更接近第一车道约束810。

[0223] 图10图示了与公开的实施方式相一致的用于导航车辆200的过程1000。过程1000可识别限定了车辆行进的行进车道的车道约束。过程1000包括：响应于检测车道偏移条件而致使车辆200行进得更接近一个或另一个车道约束，诸如第一车道约束810或第二车道约束820。以这种方式，过程1000可用于在车道约束内驾驶车辆200以及考虑可能更期望使车辆200在第一车道约束810和第二车道约束820内移位的环境因素。

[0224] 在步骤1010，过程1000可包括：使用至少一个图像捕获装置122、124和/或126来获取车辆200附近区域的多个图像。例如，处理单元110可经由数据接口128接收多个图像。例如，处理单元110可构造成确定一个或多个对象825存在于道路的一侧上。例如，对象825可以是停放或静止的车辆、墙或人(诸如行人)或自行车交通。处理单元110可基于对象825的某些特性来确定车道偏移条件是否存在。例如，处理单元110可确定：如果对象825的高度超过预定阈值，诸如距路面10cm，则对象825构成车道偏移条件。以这种方式，处理单元110可构造成过滤出将不会出现车道偏移条件的对象，诸如小件碎片。

[0225] 另外地或替代地，处理单元110可构造成基于目标对象830在车辆第一侧802上的位置来确定车道偏移条件是否存在于第一车辆侧802上。目标对象830可以是构成车道偏移条件的一个或多个对象。另外地或替代地，目标对象830可不构成车道偏移条件。例如，目标对象830可以是道路障碍或街道标志。处理单元110可构造成：如果这样的目标对象830被检测则得出结论车道偏移条件不存在。

[0226] 在步骤1020，过程1000可根据多个图像确定在第一车辆侧802的第一车道约束和在第二车辆侧812的第二车道约束。例如，处理单元110可构造成基于经由数据接口128接收的多个图像来确定第一车道约束810和第二车道约束820。

[0227] 在步骤1030，过程1000可确定车道偏移条件是否存在于第一车辆侧802上。如果车道偏移条件确实存在于第一车辆侧802上，则在步骤1040，过程1000可包括：致使车辆200在第一车道约束810和第二车道约束820内行进，使得第二距离815小于第一距离805。处理单元110可构造成致使车辆200行进，使得第一距离805和第二距离810基本上不同，以避免或

IPR2025-00943
Tesla EX1002 Page 1283

保持更远离对象825。例如，如果车道偏移条件存在于第一车辆侧802上，则处理单元110可构造成致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805大于第二距离815。这例如在图8A中图示。替代地，如果车道偏移条件存在于第二车辆侧812上，则处理单元110可构造成致使车辆200在第一车道约束810和第二车道约束820内行进，使得第一距离805小于第二距离815。这例如在图8B中图示。例如，为了根据过程1000导航车辆200，处理单元110可向车辆200的节流系统220、制动系统230和/或转向系统240中的一个或多个发送电子信号。

[0228] 如果车道偏移条件不存在于第一车辆侧802上，则在步骤1050中，方法可包括：确定车道偏移条件是否存在于第二车辆侧812上。如果车道偏移条件存在于第二车辆侧812上，所述方法可包括：致使车辆200在第一车道约束810和第二车道约束820之间行进，使得第一距离805小于第二距离815。

[0229] 将车辆导航至默认车道

[0230] 系统100可提供驾驶员辅助功能性，用于监测车辆200在当前行进车道中的地点并且如果当前行进车道并非预定的默认行进车道则使车辆200进入预定的默认行进车道。预定的默认行进车道可由用户输入限定和/或作为处理从至少一个图像捕获装置122-126接收的图像的结果。

[0231] 系统100还可为车辆200的驾驶员和乘客和/或其它驾驶员提供即将发生车辆200向预定的默认行进车道变道的通知。例如，系统100可激活定位在保险杠区域210之一处或附近的转向信号。另外地或替代地，系统100可诸如通过扬声器360发出语音通知。当车辆行进时，车辆200可使用系统100将特定车道选择为预定的默认行进车道。当车辆200行进的当前车道(例如，当前行进车道)不同于预定的默认行进车道时，系统100可致使车辆200做出变道，例如，进入预定的默认行进车道。

[0232] 图11图示了在道路800的当前行进车道1110中行进的车辆200。根据一些实施方式，车辆200可具有预定的默认行进车道1120。例如，预定的默认行进车道1120可以是道路800的最右车道。导航系统100可构造成确定车辆200是否行进在预定的默认行进车道1120中，如果否，则致使车辆200改变至预定的默认行进车道1120。例如，处理单元110可构造成将当前行进车道1110与预定的默认行进车道1120进行比较，并且致使车辆200返回到预定的默认行进车道1120。

[0233] 根据一些实施方式，处理单元110可基于由图像捕获装置122-126中的一个或多个获取的多个图像来检测预定的默认行进车道1120中的对象1130。处理单元110可通过使用上面结合图5A至图5D描述的技术来分析图像而检测车辆200相对于对象1130的位置和速度。根据一些实施方式，处理单元110可确定车辆200是否应该超越对象1130。这可能期望的是，例如车辆200相对于对象1130的速度和/或加速度超过某一阈值。

[0234] 图12A是与公开的实施方式相一致的存储器140的示例性框图，存储器140可存储用于执行一个或多个操作的指令。如图12A中图示的，存储器140可存储用于执行本文中描述的默认行进车道和当前行进车道的检测和响应的一个或多个模块。例如，存储器140可存储当前车道检测模块1210、默认车道检测模块1220和动作模块1230。

[0235] 当前车道检测模块1210可存储指令，当由处理单元110执行时，该指令可检测并且限定当前行进车道1110。例如，处理单元110可执行当前车道检测模块1212，以处理从至少

IPR2025-00943
Tesla EX1002 Page 1284

一个图像捕获装置122-124接收的多个图像并且检测当前行进车道1110。这可包括检测车辆200相对于道路800的路边、道路800的中点和/或车道标记物(诸如喷涂的车道线)的相对位置。处理单元110可通过使用上面结合图5A至图5D描述的技术来分析图像而检测车辆200相对于道路800的位置。

[0236] 默认车道检测模块1220可存储指令,当由处理单元110执行时,该指令可识别预定的默认行进车道1120并且将默认车道检测模块1220与由当前车道检测模块1210识别的当前行进车道1110进行比较。默认车道检测模块1220可基于用户输入和/或从至少一个图像捕获装置122-124接收的多个图像来识别预定的默认行进车道1120。例如,用户可通过扬声器360接收的语音命令或者通过选择显示在触摸屏320上的菜单来选择默认车道。

[0237] 在一些实施方式中,预定的默认行进车道1120可以是最右行进车道,或者可以是具有更期望特征的车道。例如,默认车道检测模块1220可包括指令,该指令用于选择交通量更少的车道、最接近车辆200速度行进的车辆的车道和/或最接近一个接近出口或车辆200将转向的街道的车道。在一些实施方式中,预定的默认行进车道1120可基于从位置传感器130(诸如全球定位系统)接收的数据来确定。

[0238] 动作模块1230可存储指令,当由处理单元110执行时,该指令可产生识别当前行进车道1110和预定的默认行进车道1120的响应。例如,处理单元110可执行动作模块1230,以在当前行进车道1110与预定的默认行进车道1120不相同的情况下致使车辆200变道。另外地或替代地,动作模块1230可启动变道的通知,诸如激活转向信号210或者产生语音通知。

[0239] 动作模块1230可包括致使处理单元110确定车辆200是否应该移出预定的默认行进车道1120的指令。例如,动作模块1230可确定:预定的默认行进车道1120中的对象1130应该由车辆200绕过。

[0240] 动作模块1230可进一步包括致使处理单元110确定进行变道是否安全的指令。例如,在执行变道之前,处理单元110可基于例如使用上面结合图5A至图5D描述的技术进行的多个图像的分析来评估确定道路800的一个或多个特性。如果分析指示变道条件是不安全的,则处理单元110可确定:车辆200应该保持在当前车道中。

[0241] 动作模块1230可包括用于操纵车辆200以机动离开和返回到预定的默认行进车道1120(诸如绕过对象1130)的额外指令。当执行这些指令时,处理单元110可使用车辆200的估算运动来估算预定的默认行进车道1120的地点。这在以下情形下可能是有利的,即预定的默认行进车道1120不能保持为由图像捕获装置122-124检测到。

[0242] 处理单元110可构造成执行偏移轮廓,该偏移轮廓使车辆200从预定的默认行进车道1120偏移完整的车道宽度,其可基于被检测的车道标记物之间的距离来估算。由偏移轮廓限定的偏移机动可由处理单元110执行,使得车辆200的移动对驾驶员而言平稳而舒适。例如,偏移摆率可大约为0.15米/秒到0.75米/秒。在偏移机动期间,处理单元110可执行来自动作模块1230的指令,该指令切换至跟随当前行进车道1110的中央。这些指令可包括从跟随带偏移的预测车道逐渐过渡到跟随当前行进车道1110。逐渐过渡可通过限定期望的前瞻点来完成,使得不允许在时间中的改变多于建议的摆率(可约为0.25米/秒)。

[0243] 用于使车辆200从预定的默认行进车道1120过渡的指令可用于使车辆200过渡回到预定的默认行进车道中。例如,处理单元110可构造成执行偏移轮廓,该偏移轮廓使车辆200从当前行进车道1110偏移完整车道宽度,其可基于被检测的车道标记物之间的距离来

IPR2025-00943
Tesla EX1002 Page 1285

估算。由偏移轮廓限定的偏移机动可由处理单元110执行，使得车辆200的移动对驾驶员而言平稳而舒适。例如，偏移摆率可大约为0.15米/秒到0.75米/秒。在偏移机动期间，处理单元110可执行来自动作模块1230的指令，该指令切换至跟随预定的默认行进车道1120的中央。这些指令可包括从跟随带偏移的预测车道逐渐过渡到跟随预定的默认行进车道1120。逐渐过渡可通过限定期望的前瞻点来完成，使得不允许在时间中的改变多于建议的摆率(可约为0.25米/秒)。

[0244] 在使车辆200在两个车道之间过渡的一个或多个步骤，处理单元110可基于前瞻点和偏移轮廓来计算并实施偏航率命令。偏航率命令可用于基于偏航率命令和使用陀螺仪测量的车辆200的偏航率来计算扭矩命令。处理单元110可向车辆200的转向系统(例如，转向系统240)发送扭矩命令。

[0245] 图12B图示了与公开的实施方式相一致的用于将车辆200导航至默认车道的过程1250。根据一些实施方式，过程1250可由导航系统100的一个或多个部件(诸如至少一个处理单元110)实施。过程1250可识别车辆200的当前行进车道1110并且将之与预定的默认行进车道1120比较，以确定车辆200是否行进在预定的默认行进车道1120内。如果车辆200并未行进在预定的默认行进车道1120中，则处理单元110可致使车辆200变道，使得导航至默认行进车道1120。

[0246] 在步骤1260，图像捕获装置122-124中的一个或多个可获取车辆200附近区域的多个图像。处理单元110可经由数据接口128接收多个图像。

[0247] 在步骤1270，处理单元110可执行当前车道检测模块1210以检测并且限定当前行进车道1110。例如，处理单元110可处理多个图像，以例如检测车辆200相对于道路100的路边、道路100的中点和/或喷涂的车道标记物的相对位置。处理单元110可通过使用上面结合图5A至图5D描述的技术来分析图像而检测车辆200相对于道路800的位置。

[0248] 在步骤1280，处理单元110可执行默认车道检测模块1220，以识别预定的默认行进车道1120并且将默认车道检测模块1220与由当前车道检测模块1210识别的当前行进车道1110比较。默认车道检测模块1220可基于用户输入和/或从至少一个图像捕获装置122-124接收的多个图像来识别预定的默认行进车道1120。

[0249] 例如，如上面讨论的，处理单元110可基于任何数目的因素来确定预定的默认行进车道1120。预定的默认行进车道1120可以是在多个行进车道1110和1120之中的最右车道。另外地或替代地，预定的默认行进车道1120的确定可基于通过处理多个图像经由用户接口170从用户接收的输入，和/或基于从位置传感器130(诸如全球定位系统)接收的数据。预定的默认行进车道1120可在每次车辆行程确定一次，或者定期确定。例如，预定的默认行进车道1120可响应于给定地点的条件动态地确定。

[0250] 如果当前车道1110与预定的默认行进车道1120不相同，则处理单元110可执行动作模块1230以致使车辆100导航至默认行进车道1120。例如，处理单元110可向车辆200的节流系统220、制动系统230和/或转向系统240中的一个或多个发送电子信号，以触发期望的响应。例如，处理单元110可致使转向系统240将车辆200的方向盘转向以实现预定角度的旋转。

[0251] 进一步，在执行变道之前，处理单元110可确定变道是否安全(例如，路中没有车辆或对象)。如上面讨论的，处理单元110可基于例如使用上面结合图5A至图5D描述的技术进

IPR2025-00943
Tesla EX1002 Page 1286

行的多个图像的分析来评估道路800的一个或多个特性。如果变道条件是不安全的，则处理单元110可确定：车辆200应该保持在当前车道中。

[0252] 在一些实施方式中，执行动作模块1230可启动变道的通知。例如，处理单元110可产生信号以激活转向信号和/或经由扬声器360产生语音通知。因此，这样的通知可发生在车辆200变道之前。在其它实施方式中，执行动作模块1230可向车辆200的驾驶员启动在不致使变道发生的情况下则车辆200并未行进在默认车道中的通知。因此，在这样的实施方式中，车辆200的驾驶员可确定是否手动采取动作以及使车辆200转向以变道。

[0253] 在一些实施方式中，系统100可通知车辆200的用户和/或车辆200附近的其它驾驶员：如果车辆200将移向预定的默认行进车道1120，则车辆200将变道。例如，处理单元110可在致使车辆200变道之前激活车辆200的转向信号。另外地或替代地，通知附近的其它驾驶员的语音通知可由处理单元110产生。语音通知可经由例如蓝牙发送到包括在其它近旁车辆中的接收器。另外地或替代地，至少一个处理单元110可在致使车辆200切换车道之前致使语音指示器通知车辆200的驾驶员。以这种方式，车辆200的驾驶员可以预料：控制系统100将施加在车辆200上。

[0254] 控制转弯车辆的速度

[0255] 系统100可提供驾驶员辅助功能性，用于控制车辆200在不同场景(诸如做出转弯(例如，通过弯道)时)下的速度(即，速度和/或方向)。例如，系统100可构造成使用地图、位置和/或视觉数据的组合，以控制车辆200在转弯时的速度。特别是，系统100可控制车辆200在接近弯道时(在导航弯道的同时)和/或在离开弯道时的速度。系统100可考虑用以控制速度的不同因素，这取决于在给定的时间车辆200相对于弯道的位置。例如，系统100可初始响应于获知到弯道位于车辆200之前的某一距离来致使车辆200减速。然而，在车辆200转弯经过弯道的同时，系统100可基于对弯道特性的分析来致使车辆200加速和/或改变方向。

[0256] 图13A和图13B是与公开的实施方式相一致的使车辆(例如，车辆200)接近并导航具有道路1300的一个或多个特性的弯道的概略表示。如图13A中图示的，视觉指示器(诸如位于道路1300附近的道路标志1302)可告知车辆200弯道在前面。响应于由道路标志1302提供的警告，车辆200的驾驶员和/或为车辆200提供驾驶员辅助功能性的系统100可调整车辆200的速度(例如，通过减速和/或转向)，以安全地导航弯道。

[0257] 在车辆200接近弯道时，系统100可通过使用车辆200的图像捕获装置122和/或124(可以是相机)来检测车道宽度1306和弯道半径1308。基于弯道半径1308、车道宽度1306以及使用图像捕获装置122和/或124检测的与弯道关联的其它特性，系统100可调整车辆200的速度。具体地，系统100可调整车辆200接近弯道和/或车辆200导航弯道时的速度。系统100还可响应于例如出现在弯道之后的交通灯1304来调整车辆200离开弯道时的速度。下面结合图14和图15描述关于调整车辆200的速度以及检测车道宽度1306、曲率半径或弯道半径1308及与弯道关联的特性的进一步细节。

[0258] 在一些实施方式中，系统100可基于地图数据和/或车辆位置信息(例如，GPS数据)来辨识将被导航的弯道。系统100可基于反映在地图数据中的弯道的一个或多个特性来确定车辆的初始目标速度。系统100可将车辆的速度调整至初始目标速度，并且基于由图像捕获装置122-126中的一个或多个获取的一个或多个图像来确定弯道的观察到的一个或多个特性。系统100可基于弯道的观察到的一个或多个特性来确定更新的目标速度，并且将车辆

IPR2025-00943
Tesla EX1002 Page 1287

的速度调整至更新的目标速度。

[0259] 图14是与公开的实施方式相一致的存储器的示例性框图，该存储器被构造成存储用于执行一个或多个操作的指令。如图14中指示的，存储器140可存储弯道辨识模块1402、特性观察模块1404和速度模块1406。公开的实施方式并不限于任何特定构造的存储器140。进一步，处理单元110可执行存储在任何模块1402-1406(包括在存储器140中)中的指令。

[0260] 在一个实施方式中，弯道辨识模块1402可存储软件指令(诸如计算机视觉软件)，当由处理单元110执行时，该软件指令通过例如分析一组或多组图像和/或使用地图和/或位置数据(例如，存储在地图数据库160中的数据)来检测弯道。图像组可由图像捕获装置122、124和126中的一个或多个获取。基于该分析，系统100(例如，经由处理单元110)可致使车辆200的速度改变。例如，如果处理单元110检测到存在位于车辆200之前的最小阈值距离内的弯道，则处理单元110可致使车辆200的速度降低预定的量和/或将车辆200的方向调整特定的角度。

[0261] 在一个实施方式中，特性观察模块1404可存储软件指令(诸如计算机视觉软件)，当由处理单元110执行时，该软件指令观察与系统100(例如，经由执行弯道辨识模块1402)检测的弯道关联的特性。处理单元110可通过分析由图像捕获装置122、124和126中的一个或多个获取的一组或多组图像和/或使用地图和/或位置数据(例如，存储在地图数据库160中的数据)来观察与被检测弯道关联的特性。该分析可产生与弯道关联的信息，诸如弯道半径1308、车道宽度1306、曲率的程度、曲率变化率、转弯倾斜度、弯道的长度或弧长，等。这些特性可包括估值(例如，使用预先存在的数据和/或基于这样的数据的数学模型计算的弯道特性)和实际值(例如，通过分析弯道的被捕获图像计算的弯道特性)。基于该分析，处理单元110可致使车辆200的速度改变。例如，鉴于与位于车辆200之前25米的弯道关联的曲率变化率较高，处理单元110可致使车辆200的速度降低和/或将车辆200的方向调整特定的角度。

[0262] 在一个实施方式中，速度模块1406可存储被构造成分析来自一个或多个计算机电装置的数据的软件指令，该一个或多个计算机电装置被构造成确定目标速度并且致使车辆200的速度改变为目标速度。例如，处理单元110可执行速度模块1406，以基于从执行弯道辨识模块1402和特性观察模块1404得到的数据来计算车辆200的目标速度。例如，这样的数据可包括初始目标位置和初始速度、更新目标位置和更新速度、最终目标位置和最终速度，等。另外，处理单元110可基于来自车辆200的其它系统(诸如车辆200的节流系统220、制动系统230和/或转向系统240)的输入来计算车辆200的速度。基于计算的目标速度，处理单元110可向车辆200的节流系统220、制动系统230和/或转向系统240发送电子信号(例如，经由控制器区域网络总线(“CAN总线”))，以通过例如物理踩下制动器、松开车辆200的加速器或者使车辆200沿特定方向转向(例如，以特定角度)来触发车辆200的速度和/或方向的改变。速度增加可与对应的加速度(例如，介于0.2米/秒^{2}到3.0米/秒^{2}之间)关联。反过来，速度降低可与对应的减速度关联。另外，加速度或减速度可基于道路类型(例如，高速公路、城市街道、乡村道路，等等)、近旁任何速度约束的存在(例如，急剧弯道、交通灯)以及道路宽度或车道宽度1306。出于本公开的目的，减速可指具有负值的加速度。

[0263] 图15是示出与公开的实施方式相一致的用于基于被检测弯道和观察的弯道特性和/或关于弯道的地图数据来控制车辆速度的示例性过程1500的流程图。在步骤1510，处理

IPR2025-00943
Tesla EX1002 Page 1288

单元110可经由位于处理单元110与图像获取单元120之间的数据接口128接收多个图像。例如，包括在图像获取单元120中的相机(诸如图像捕获装置122)可捕获多个图像并且经数字连接(例如，USB、无线、蓝牙，等等)将图像发送到处理单元110。在一些实施方式中，处理单元110可经由多个数据接口接收多于一个的多个图像。例如，处理单元110可从图像捕获装置122、124、126中的每个接收多个图像，图像捕获装置122、124、126中的每个均可具有关联的数据接口用于与处理单元110通信数据。公开的实施方式并不限于任何特定的数据接口构造或协议。

[0264] 在步骤1520，处理单元110可执行弯道辨识模块1502，以基于例如地图和/或位置数据(例如，存储在地图数据库160中的数据)来检测位于车辆200之前的弯道。这样的数据可指示存在于于给定道路上的弯道的地点。数据可指示特定弯道的绝对位置(例如，经由GPS坐标)或弯道的相对位置(例如，通过相对于车辆200、道路标志1302、交通灯1304和/或其它描述性地标来描述弯道)。在一些实施方式中，处理单元110可执行弯道辨识模块1502，以通过分析多个图像来检测弯道。该分析可使用上面结合图5A至图5D和图6描述的技术来进行。

[0265] 在步骤1530，处理单元110可执行速度模块1406，以基于例如从执行弯道辨识模块1402得到的数据来确定车辆200的初始目标速度。初始目标速度可反映车辆200的速度和方向以安全地进入即将到来的弯道，并且可使用上面结合图14描述的技术来确定。例如，速度模块1406可基于与即将到来的弯道关联的已知或估算的特性(诸如弯道半径1308和转弯倾斜度)来确定初始目标速度。速度模块1406可依据数学模型和/或公式并且基于各种约束来计算初始目标速度。在一些实施方式中，例如，速度模块1406可基于交通灯(诸如交通灯1304)存在于弯道之前、期间或之后、天气条件(例如，小雨、大雨、大风、多云，等等)、道路条件(例如，柏油路、土路、坑洼，等等)、限速、近旁存在车辆或行人等等来确定初始目标速度。初始目标速度还可基于与车辆200关联的侧向加速极限。在一些实施方式中，侧向加速极限可在0.2米/秒^{2}到3.0米/秒^{2}的范围内并且可取决于弯道特性、道路宽度或车道宽度1306来调整。另外，初始目标速度可基于与一个或多个相机(诸如图像捕获装置122和124)关联的焦距范围。

[0266] 例如，速度模块1406可基于车辆200之前的道路的曲率半径来确定初始目标速度。速度模块1406可确定距车辆200的各种距离处的曲率半径并且以向量R(1:n)存储信息。在一个实施方式中，速度模块1406可确定以下地点的曲率半径：(1)车辆200的当前位置；(2)距车辆200的距离$D_{终点}$，其中$D_{终点}$表示车辆200之前的、道路特性的认识允许系统100控制车辆200的导航响应使得驾驶员可感觉舒适的距离；以及(3)距车辆200的增量距离，诸如每1米，等等。距离$D_{终点}$的值可取决于车辆200的当前速度V_{ego}。例如，V_{ego}的值越大(例如，车辆200行进得越快)，$D_{终点}$的值越大(例如，可考虑越长的距离为系统100提供足够的时间来提供导航响应，诸如基于接近急剧转弯或弯道来舒适地减小车辆200的速度)。在一个实施方式中，$D_{终点}$可根据以下公式计算：$D_{终点}=(V_{ego}^2/2)/(0.7/40*V_{ego}+0.5)$。由此，向量D(1:n)＝[1:1:n]，其中n＝向下取整($D_{终点}$)。

[0267] 在一些实施方式中，速度模块1406可基于侧向加速度约束$a_{最大}$(例如，最大允许的侧向加速度)来确定初始目标速度。$a_{最大}$可落入1.5米/秒^{2}到3.5米/秒^{2}的范围内。基于$a_{最大}$和(上面讨论的)曲率半径向量R，速度模块1406可将(例如，距车辆200的当前位置的不同点或

IPR2025-00943
Tesla EX1002 Page 1289

距离处的)车辆200的速度限制为速度$v_{极限}$,计算为:$v_{极限}(1:n)=(a_{最大}*R(1:n))^{1/2}$。另外,根据$v_{极限}$的确定而导出的减速度约束可计算为:$dec_{极限}$=最小{$((v_{极限}(1:n)^2-V_{ego}^2)/2)$,最大{$d_{最小}$,$D(1:n)-2*V_{ego}$}},其中$d_{最小}$等于40米。$d_{最小}$的值可表示对车辆200的驾驶员而言减速感觉舒适的点。

[0268] 在一些实施方式中,速度模块1406可基于曲率摆率约束$slew_{最大}$(可基于曲率变化率)来确定初始目标速度。$slew_{最大}$可落入0.5米/秒^{3}到1.5米/秒^{3}的范围内。速度模块1406可根据以下公式来确定$v_{极限}$:$v_{极限}(1:n)=(slew_{最大}/abs\{\{1/R(1:n)\}'\})^{1/3}$。进一步,速度模块1406可根据以下公式来确定$dec_{极限}$:$dec_{极限}$=最小{$((v_{极限}(1:n)^2-V_{ego}^2)/2)$,最大{$d_{最小}$,$D(1:n)-2*V_{ego}$}},其中$d_{最小}$等于40米。

[0269] 在一些实施方式中,速度模块1406可确定舒适约束,使得在车辆200经历各种导航响应(诸如转弯或导航弯道)的同时车辆200的驾驶员是舒适的。速度模块1406可根据以下公式来确定$v_{极限}$:$v_{极限}(1:n)=4/500*R(1:n)+22.5$。进一步,速度模块1406可根据以下公式来确定$dec_{极限}$:$dec_{极限}$=最小{$((v_{极限}(1:n)^2-V_{ego}^2)/2)$,最大{$d_{最小}$,$D(1:n)-2*V_{ego}$}},其中$d_{最小}$等于40米。

[0270] 速度模块1406可通过使用与每个约束关联计算的$dec_{极限}$的最小值来合并上面确定的三个约束,即侧向加速度约束、曲率摆率约束和舒适约束。在合并约束之后,速度模块1406可根据以下公式通过确定调整车辆200的速度的量来达到初始目标速度,记作$v_{舒适}$:$v_{舒适}=v_{舒适(先前)}+dec_{极限}*\Delta t$,其中$v_{舒适(先前)}$表示车辆200的先前速度(例如,先前确定的速度、在速度模块1406做任何调整之前的车辆200的速度,等等),并且Δt表示确定$v_{舒适}$与$v_{舒适(先前)}$之间的时间。

[0271] 在步骤1540,处理单元110可执行速度模块1406以将车辆200的速度调整至在步骤1530确定的初始目标速度。基于车辆200的当前速度以及在步骤1530确定的初始目标速度,处理单元110可向车辆200的节流系统220、制动系统230和/或转向系统240发送电子信号,以将车辆200的速度改变至初始目标速度。进一步,一个或多个致动器可控制节流系统220、制动系统230和/或转向系统240。例如,处理单元110可发送电子信号,致使系统100物理地踩下制动器达预定的量或者部分地松开车辆200的加速器。进一步,处理单元110可发送电子信号,致使系统100将车辆200沿特定方向转向。在一些实施方式中,加速度(或减速度)的幅值可基于车辆200的当前速度与初始目标速度之间的差。当该差较大时(例如,当前速度比初始目标速度大10km/hr),例如,系统100可通过采取使减速度最大化并且导致车辆200在最短的时间量内尽可能安全地实现初始目标速度的方式制动来改变当前速度。替代地,系统100可采取使减速度降到最低并且导致车辆200逐渐实现初始目标速度的方式来施加制动。更一般地,系统100可将当前速度改变至初始目标速度,根据任何特定的制动轮廓(例如,前两秒要求高制动水平且随后三秒要求低制动水平的制动轮廓)。公开的实施方式并不限于任何特定的制动轮廓或制动方式。

[0272] 在步骤1550,处理单元110可执行特性观察模块1404,以基于例如分析使用上面结合图5A至图5D描述的技术进行的多个图像来确定位于车辆200之前的弯道的一个或多个特性。例如,所述特性可包括与下列关联的估值和/或实际值:弯道半径1308、道路宽度或车道宽度1306、曲率的程度、曲率变化率、转弯倾斜度,等。例如,所述特性还可包括由一个或多个道路标志(诸如道路标志1302)传达的信息。此外,所述特性可包括被检测弯道的长度或

IPR2025-00943
Tesla EX1002 Page 1290

弧长(可以是车辆200在接近弯道时的速度的函数)。处理单元110可执行特性观察模块1404,以基于数学模型和/或公式来计算与这些特性(例如,弯道半径1308、曲率的程度、曲率变化率、弯道弧长,等等)关联的值。在一些实施方式中,处理单元110可执行特性观察模块1404,以构建将弯道曲率表达为距离(例如,行进的距离)的函数的数学模型。这样的模型可能受约束,诸如最大侧向加速度约束(例如,在1.5米/秒^{2}到3.0米/秒^{2}的范围内)和/或最大侧向加速度衍生约束(例如,在0.8米/秒^{2}到1.2米/秒^{2}的范围内)。这些约束可产生车辆200的最大速度(作为行进距离的函数)。

[0273] 在步骤1560,处理单元110可执行速度模块1406,以基于在步骤1550确定的特性来计算更新的目标速度。更新的目标速度可反映车辆200的更新速度和/或方向以安全地执行下列的任何组合:(i)进入即将到来的弯道;(ii)导航通过弯道;或(iii)离开弯道。更新的目标速度可使用上面结合图14和步骤1530描述的技术来确定。例如,速度模块1406可基于与弯道关联的特性并且基于数学模型和/或公式来确定更新的目标速度。

[0274] 在步骤1570,处理单元110可执行速度模块1406,以将车辆200的速度调整至在步骤1560确定的更新的目标速度。这种调整可使用上面结合步骤1540描述的技术来完成。例如,基于车辆200的当前速度以及在步骤1560确定的更新的目标速度,处理单元110可向车辆200的节流系统220、制动系统230和/或转向系统240发送电子信号(例如,经由CAN总线),以将车辆200的速度改变至更新的目标速度。在这一点,车辆200的当前速度不一定与在步骤1530确定的初始目标速度相同。例如,处理单元110可在步骤1560已经确定更新的目标速度,同时将车辆200的速度调整至初始目标速度。这可能真正存在于车辆200的速度与初始目标速度的差较大(例如,大于10km/hr)并且关联的加速度较小的情况下。

[0275] 在一些实施方式中,系统100可基于对弯道特性的继续观察(例如,基于描绘弯道的车辆200接近的部分的图像数据)来常规更新到目标速度。另外,在系统100将车辆200导航通过弯道之后,系统100可致使车辆200加速至新的目标速度,该新的目标速度适于行进道路的不带弯道的区段,这通过分析由系统100获取的地图数据、位置数据和/或图像数据而确定。

[0276] 模仿前车

[0277] 系统100可提供驾驶员辅助功能性,用于致使车辆200模仿(或者拒绝模仿)不同场景(诸如当前车切换车道、加速或转弯时)下的前车。例如,系统100可构造成通过分析多个图像来检测前车并且确定前车采取的一个或多个动作。在一些场景下,诸如当前车在路口处转弯时,系统100可构造成致使车辆200拒绝模仿转弯。在其它场景下,诸如当前车在路口处转弯且在路口内不变道时,系统100可构造成致使车辆200模仿转弯。系统100还可构造成基于前车的导航历史致使车辆200模仿前车的一个或多个动作。

[0278] 图16A和图16B是与公开的实施方式相一致的模仿前车200b在道路1600a上的一个或多个动作的主车辆200a的概略表示。如图16A和图16B中图示的,主车辆200a可在道路1600a的同一车道中行进的同时追踪前车200b。车道可具有左边缘1620、右边缘1630和中点1610,w_1和w_2分别指示车道的左半边和右半边。主车辆200a和前车200b可定位在车道中,使得车辆都不居中于中点1610上。如图所示16A,例如,主车辆200a可居中于中点1610的右边,使得主车辆200a与左边缘1620之间的距离c_1大于主车辆200a与右边缘1630之间的距离c_2。对比之下,前车200b可居中于中点1610的左边,使得前车200b与左边缘1620之间的距离c_3

IPR2025-00943
Tesla EX1002 Page 1291

小于前车200b与右边缘1630之间的距离c_4。主车辆200a和前车200b的位置可以有不同的构造。例如，主车辆200a和前车200b可定位在道路1600a上的任何地方。进一步，虽然图16A描绘了具有相机系统的前车200b，但是公开的实施方式并不限于前车200b包括相机系统的构造。

[0279] 主车辆200a的系统100可致使主车辆200a模仿前车200b的一个或多个动作。如图16B所示，例如，主车辆200a可模仿前车200b执行的车道移位。在模仿车道移位之后，距离c_1可等于距离c_3，并且距离c_2可等于距离c_4。在一些实施方式中，主车辆200a可模仿前车200b执行的车道移位，使得距离c_1和c_3与c_2和c_4不等。例如，在主车辆200a模仿主车辆200b的向左车道移位并且(在车道移位之后)定位至主车辆200a的左边的情况下，c_1将小于c_3并且c_2将大于c_4。替代地，在主车辆200a模仿主车辆200b的向左车道移位并且定位至主车辆200a的右边的情况下，c_1将大于c_3并且c_2将小于c_4。

[0280] 图16C和图16D是与公开的实施方式相一致的模仿前车200b在道路1600b上的一个或多个动作的主车辆200a的概略表示。如图16C中图示的，主车辆200a可在前车200b后面行进，前车200b可在路口处左转弯。主车辆200a的系统100可致使主车辆200a模仿前车200b的左转弯，如图16D所示。系统100可基于各种因素(诸如前车200b是否在路口内转弯的同时变道、与前车200b关联的导航历史、主车辆200a与前车200b之间的相对速度差，等)来致使主车辆200a模仿左转弯。另外，主车辆200a可在各种条件下(诸如基于前车200b是否沿着与主车辆200a行进沿着的路线的相同路线或相同部分行进)模仿前车200b。关于系统100可致使主车辆200a模仿前车200b的一个或多个动作的场景的进一步细节在下面结合图17至图22描述。

[0281] 在一些实施方式中，系统100可构造成确定是否模仿前车的特定移动。例如，系统100可确定：前车的某些移动不需要被模仿，因为这种移动不影响前车的路线。例如，系统100可确定：移动的少量改变(诸如在特定车道内移动)不需要被模仿。在其它实施方式中，系统100可实施平稳操作以过滤出前车的少量移动。平稳操作可致使主车辆实施前车的更显著和/或重要的移动(例如，变道)，同时过滤出较少量的不会影响前车的整体路线的移动。

[0282] 图17是与公开的实施方式相一致的存储器的示例性框图，该存储器被构造成存储用于执行一个或多个操作的指令。如图17所示，存储器140可存储车道约束模块1702、动作检测模块1704和动作响应模块1706。公开的实施方式并不限于任何特定构造的存储器140。进一步，处理单元110可执行存储在任何模块1702-1706(包括在存储器140中)中的指令。

[0283] 在一个实施方式中，车道约束模块1702可存储软件指令(诸如计算机视觉软件)，当由处理单元110执行时，该软件指令确定与主车辆200a的行进车道关联的一个或多个约束。车道约束可包括车道的中点1610、左边缘1620和右边缘1630。车道约束还可包括标记在道路1600a或1600b的表面上的一个或多个线或其它符号。处理单元110可执行车道约束模块1702，以通过例如分析一组或多组图像和/或使用地图和/或位置数据(例如，存储在地图数据库160中的数据)来确定约束。图像组可由图像捕获装置122、124和126中的一个或多个获取。基于该分析，系统100(例如，经由处理单元110)可致使主车辆200a在由左边缘1620和右边缘1630限定的车道内行进。在一些实施方式中，如果距离c_1或c_2之一小于预定阈值，则系统100可致使主车辆200a远离左边缘1620或右边缘1630。

IPR2025-00943
Tesla EX1002 Page 1292

[0284] 在一个实施方式中，动作检测模块1704可存储软件指令(诸如计算机视觉软件)，当由处理单元110执行时，该软件指令检测在主车辆200a前方行进的前车200b采取的一个或多个动作。前车200b可在与主车辆200a相同或不同的车道中行进。处理单元110可执行动作检测模块1704，以通过例如分析由图像捕获装置122、124和126中的一个或多个获取的一组或多组图像来检测前车200b采取的一个或多个动作。基于该分析，例如，处理单元110可确定前车200b已变换车道、已转弯、加速、减速、施加制动等。作为另一示例，前车200b可执行机动，使得前车200b的邻近第一车道约束的那一侧的第一偏移距离不同于前车200b的邻近第二车道约束的那一侧的第二偏移距离。例如，处理单元110可基于上面结合图5A至图5D描述的技术来执行这种分析。

[0285] 在一个实施方式中，动作响应模块1706可存储软件指令(诸如计算机视觉软件)，当由处理单元110执行时，该软件指令基于系统100(例如，经由执行动作检测模块1704)检测的前车200b所采取的一个或多个动作来确定主车辆200a采取的一个或多个动作。例如，处理单元110可执行动作响应模块1706，以确定是否模仿前车200b的一个或多个动作。该确定可基于前车200b的动作(例如，转弯、车道移位，等等)的性质、与主车辆200a关联的信息(例如，距车道路边的速度、距离c_1和c_2，等等)、道路和环境条件(例如，坑洼、大雨或风，等等)，等。在处理单元110确定模仿前车200b的一个或多个动作的情况下，处理单元110可通过以下手段完成这：例如向主车辆200a的节流系统220、制动系统230和/或转向系统240发送电子信号(例如，经由CAN总线)，以触发转弯、车道移位、主车辆200a的速度改变和/或方向改变。处理单元110可使用上面结合图4至图7描述的技术以致使主车辆200a模仿一个或多个动作。为了模仿一个或多个动作，系统100可向节流系统220、制动系统230和转向系统240中的一个或多个提供控制信号，以(例如，通过致使加速、转弯、车道移位，等等)导航车辆200。进一步，一个或多个致动器可控制节流系统220、制动系统230和/或转向系统240。

[0286] 图18是示出与公开的实施方式相一致的用于致使主车辆200a模仿前车200b的一个或多个动作的示例性过程的流程图。在步骤1810，处理单元110可经由位于处理单元110与图像获取单元120之间的数据接口128接收多个图像。例如，包括在图像获取单元120中的相机(诸如图像捕获装置122)可捕获多个图像并且经数字连接(例如，USB、无线、蓝牙，等等)将图像发送到处理单元110。在一些实施方式中，处理单元110可经由多个数据接口接收多于一个的多个图像。例如，处理单元110可从图像捕获装置122、124、126中的每个接收多个图像，图像捕获装置122、124、126中的每个均可具有关联的数据接口用于与处理单元110通信数据。公开的实施方式并不限于任何特定的数据接口构造或协议。

[0287] 在步骤1820和1830，处理单元110可执行车道约束模块1702，为主车辆200a的每侧确定与主车辆200a的行进车道关联的约束。处理单元110可通过分析在步骤1810接收的图像来确定约束。由此，在步骤1820，处理单元110可确定在主车辆200a的第一侧上的第一车道约束(例如，主车辆200a的左侧与车道的左边缘1620之间的距离，c_1)，以及在步骤1830，确定在与第一侧相反的第二侧上的第二车道约束(例如，主车辆200a的右侧与车道的右边缘1620之间的距离，c_2)。在一些实施方式中，处理单元110可通过使用地图和/或指示主车辆200a相对于车道的中点1610、左边缘1620和右边缘1630的位置的位置数据(例如，存储在地图数据库160中的数据)来确定约束。处理单元110可基于分析在步骤1810接收的图像以及地图和/或位置数据两者来确定约束；这样做可增加与约束关联的置信水平。

IPR2025-00943
Tesla EX1002 Page 1293

[0288] 在步骤1840，处理单元110可执行车道约束模块1702，以致使主车辆200a在步骤1820和1830确定的车道约束内行进。例如，主车辆200a的左侧可定位至左边缘1620的左边，或者主车辆200a的右侧可定位至右边缘1630的右边。在这样的情景下，处理单元110可向主车辆200a的节流系统220、制动系统230和/或转向系统240发送电子信号(例如，经由CAN总线)，以致使主车辆200a调整其位置并且在车道约束内行进。

[0289] 在步骤1850，处理单元110可执行动作检测模块1704，以通过分析图像将前车200b定位在在步骤1810接收的图像内。处理单元110可使用上面结合图5A和图5B描述的技术来检测车辆在图像中的存在。对于每个被检测车辆，处理单元110可构建被检测车辆相对于主车辆200a的位置、速度(例如，速度和方向)和加速度的多帧模型。处理单元110可使用上面结合图5A和图5B描述的技术来构建这样的模型。另外，处理单元110可确定与每个被检测车辆关联的行进车道，基于例如车道或其它道路标记以及在步骤1820和1830确定的车道约束。处理单元110可确定前车200b为最接近的被检测车辆，并且在与主车辆200a相同的车道中行进。在步骤1860，处理单元110可执行动作检测模块1704，以确定步骤1850在图像内检测的前车200b所采取的一个或多个动作。例如，处理单元110可确定：前车200b已变换车道、已转弯、加速、减速，等。另外，处理单元110可分析图像，以确定与前车200b关联的特性，诸如其(例如，相对于主车辆200a的)速度和位置。这些确定可基于上面结合图5A和图5B描述的技术。在一些实施方式中，处理单元110可分析图像，以确定与前车200b执行的动作关联的特性。例如，在前车200b转弯的情况下，处理单元110可确定前车200b正好在开始转弯之前的位置信息、前车200b在完成转弯时的位置信息、转弯半径以及描述前车200b在转弯时的速度改变(例如，经过一段时间)的速度轮廓。

[0290] 在步骤1870，处理单元110可执行动作响应模块1706，以致使主车辆200a模仿在步骤1850确定的前车200b的一个或多个动作。例如，处理单元110可向主车辆200a的节流系统220、制动系统230和/或转向系统240发送电子信号(例如，经由CAN总线)，以触发主车辆200a的转弯、车道移位、速度改变和/或方向改变。另外，处理单元110可构建一组路径点，以引导主车辆200a的位置和速度改变。路径点可基于由多项式(例如，第三多项式)和与目标车辆(诸如前车200b)关联的目标坐标表示的车道标线(例如，中点1610、左边缘1620、右边缘1630，等等)。处理单元110可调整多项式的系数，使得路径点超越目标坐标。处理单元110还可基于预定偏移量使路径点(例如，相对于左边缘1620和右边缘1630)偏移。在处理单元110调整多项式系数并使路径点偏移时，处理单元110还可针对形成在路径点之间的每个得到的区段来计算局部曲率信息。

[0291] 在一些实施方式中，处理单元110(例如，经由执行动作响应模块1706)可在与两个车辆关联的某一位置、速度和/或加速度约束内致使主车辆200a模仿前车200b。例如，约束可能需要主车辆200a与前车200b之间的预定最小距离(例如，安全距离)或车辆之间的速度和/或加速度的预定最大差。由此，如果模仿前车200b执行的动作将违反一个或多个约束(例如，模仿加速将导致主车辆200a太接近前车200b)，则处理单元110可致使主车辆200a拒绝模仿前车200b的一个或多个动作。另外，甚至在模仿前车200b执行的动作(例如，车道移位)的同时，处理单元110可致使主车辆200a跟随前车200b的速度轮廓并且维持与前车200b的预定距离或安全距离。处理单元110完成这可通过致使主车辆200a：(i)当间距大于预定最小距离时，不断缩短与前车200b的距离差距；(ii)不断缩短主车辆200a与前车200b的速

IPR2025-00943
Tesla EX1002 Page 1294

度之间的差距;以及(iii)将前车200b的加速度匹配至与预定约束一致的程度。为了控制这些动作应该执行到什么程度,处理单元110可基于与前车200b的当前预测距离实时将每个动作与权重关联。由此,例如,在与前车200b的距离减小时,更大的权重可加到减速上,而非匹配前车200b的加速。

[0292] 作为另一示例,在主车辆200a以大于前车200b速度的速度接近前车200b的情况下,处理单元110可致使主车辆200a逐渐减速,使得当两个车辆之间的距离等于预定安全距离(诸如最小跟随距离)时,两个车辆的速度相同。在另一示例中,在前车200b突然出现在主车辆200a前方的情况下,致使两个车辆之间的距离小于最小跟随距离,处理单元110可致使主车辆200a逐渐减速以避免更接近前车200b,同时增加车辆之间的距离直到该距离到达最小跟随距离。

[0293] 图19是示出与公开的实施方式相一致的用于致使主车辆200a拒绝模仿前车200b的转弯的示例性过程1900的流程图。在步骤1910,如上面结合图18的步骤1810描述的,处理单元110可经由位于处理单元110与图像获取单元120之间的数据接口128接收多个图像。在步骤1920,如上面结合图18的步骤1850描述的,处理单元110可执行动作检测模块1704,以通过分析图像将前车200b定位在步骤1910接收的图像内。在步骤1925,处理单元110可执行动作检测模块1704,以使用上面结合图18的步骤1860描述的技术来确定前车200b是否在路口处转弯。如果确定前车200b在路口处转弯(步骤1925,是),则在步骤1930,处理单元110可致使主车辆200a拒绝模仿前车200b的转弯。否则(步骤1925,否),过程1900结束。

[0294] 图20是示出与公开的实施方式相一致的用于致使主车辆200a模仿或拒绝前车200b的转弯的示例性过程2000的流程图。在步骤2010,如上面结合图18的步骤1810描述的,处理单元110可经由位于处理单元110与图像获取单元120之间的数据接口128接收多个图像。在步骤2020,如上面结合图18的步骤1850描述的,处理单元110可执行动作检测模块1704,以通过分析图像将前车200b定位在步骤2010接收的图像内。在步骤2025,处理单元110可执行动作检测模块1704,以使用上面结合图18的步骤1860描述的技术来确定前车200b是否在路口处转弯。如果确定前车200b在路口处转弯(步骤2025,是),则在步骤2035,处理单元110可执行动作检测模块1704,以使用上面结合图18的步骤1860描述的技术来确定前车200b是否在路口内变道。否则(步骤2025,否),过程2000结束。如果确定前车200b在路口内变道(步骤2035,是),则在步骤2050,处理单元110可致使主车辆200a拒绝模仿前车200b的转弯。否则(步骤2035,否),在步骤2040,处理单元110可致使主车辆200a模仿前车200b的转弯。

[0295] 图21是示出与公开的实施方式相一致的用于致使主车辆200a模仿第一前车的一个或多个动作的另一示例性过程2100的流程图。在步骤2110,如上面结合图18步骤1810描述的,处理单元110可经由位于处理单元110与图像获取单元120之间的数据接口128接收多个图像。在步骤2120和2130,如上面结合图18的步骤1850描述的,处理单元110可执行动作检测模块1704,以通过分析图像将第一和第二前车定位在步骤2110接收的图像内。在步骤2135,处理单元110可执行动作检测模块1704,以使用上面结合图18的步骤1860描述的技术来确定第一前车路径的转弯半径是否小于第二前车路径。如果确定第一前车路径的转弯半径小于第二前车路径(步骤2135,是),则在步骤2140,处理单元110可致使主车辆200a模仿第一前车的转弯。否则(步骤2135,否),过程2100结束。

IPR2025-00943
Tesla EX1002 Page 1295

[0296] 图22是示出与公开的实施方式相一致的用于基于导航历史来致使主车辆200a模仿前车200b的一个或多个动作的示例性过程2200的流程图。在步骤2210，如上面结合图18的步骤1810描述的，处理单元110可经由位于处理单元110与图像获取单元120之间的数据接口128接收多个图像。在步骤2220，如上面结合图18的步骤1850描述的，处理单元110可执行动作检测模块1704，以通过分析图像将前车200b定位在步骤2210接收的图像内。

[0297] 在步骤2230，处理单元110可通过分析在步骤2210接收的图像来确定前车200b的位置信息。前车200b的位置信息可指示前车200b相对于主车辆200a的位置。处理单元110可通过使用上面结合图5A至图5D描述的技术来分析图像以得到前车200b的位置信息以及前车200b采取的一个或多个动作。

[0298] 在步骤2240，处理单元110可将在步骤2230确定的前车200b的位置信息与预定地图数据(例如，存储在地图数据库160中的数据)比较。例如，通过将位置信息与预定地图数据比较，处理单元110可得到前车200b的位置坐标。另外，处理单元110可使用预定地图数据，作为与在步骤2230确定的位置信息关联的置信水平的指示器。

[0299] 在步骤2250，处理单元110可通过例如跟踪前车200b采取的动作来创建与前车200b关联的导航历史。处理单元110可将前车200b采取的动作(例如，在步骤2230根据分析图像而确定)与步骤2240的预定地图数据关联。由此，导航历史可指示执行某些动作(例如，转弯、车道移位，等等)的前车200b处于某一位置坐标。

[0300] 在步骤2260，处理单元110可执行动作响应模块1706，以基于在步骤2250创建的导航历史并且使用上面结合图18的步骤1870描述的技术来确定是否致使主车辆200a模仿前车200b的一个或多个动作。例如，甚至在根据分析图像(例如，在步骤2230)没有检测到转弯的情况下，当导航历史指示前车200b已成功导航至少一先个前转弯时，处理单元110(例如，经由动作响应模块1706)可致使主车辆200a模仿前车200b的转弯。更一般地，在导航历史指示前车200b已成功执行这样的动作的情况下，处理单元110可致使主车辆200a模仿前车200b采取的动作。

[0301] 导航车辆以超越另一车辆

[0302] 系统100可提供驾驶员辅助功能性，用于监测车辆200附近并且倘若变道被认为置于致使诸如与另一车辆(例如，目标车辆)碰撞的风险下则中止变道。例如，如果确定目标车辆位于与车辆200不同的车道中，则系统100可使车辆200能够完成对目标车辆的超越。如果，在完成超越之前，系统100确定目标车辆正进入车辆200的行进车道，系统100可致使车辆200中止超越。例如，中止超越可包括致使车辆200制动或致使车辆200停止加速。

[0303] 如图23中图示的，车辆200可行进在道路2300上。车辆200可配备有系统100并且可实施任何公开的实施方式来识别车道约束以及在车道约束内操作车辆200。系统100还可促进车辆200做出的变道。例如，系统100可确定情形是否允许车辆200安全地变道，如果在变道期间情形改变，则致使车辆200中止变道。

[0304] 如图23所示并且如先前相对于图8A至图8D讨论的，车辆200可包括第一车辆侧802，第一车辆侧802距第一车道约束2302可以是第一距离805。类似地，车辆200可包括与第一车辆侧802相反的第二车辆侧812，并且第二车辆侧812距第二车道约束2304可以是第二距离815。以这种方式，第一车道约束810和第二车道约束2304可限定车辆200可行进的车道2330。

IPR2025-00943
Tesla EX1002 Page 1296

[0305] 系统100可基于由图像捕获装置122-126中的一个或多个获取的多个图像来确定第一车道约束2302和第二车道约束2304。根据一些实施方式，第一车道约束2302和/或第二车道约束2304可由可见的车道边界(包括标记在路面上的线)识别。另外地或替代地，第一车道约束2302和/或第二车道约束2304可包括路面的路边。根据一些实施方式，系统100可通过识别路面宽度2320的中点(或者车辆200可用作导航基准的任何其它合适的道路特征)来确定第一车道约束2302和/或第二车道约束2304。例如，当指定道路车道的线未被喷涂或以其它方式贴标签时，系统100可采取这种方式识别车道约束。

[0306] 当超越目标车辆2310时，车辆200的变道可能由于目标车辆2330的轨迹改变而复杂化。由此，系统100可基于对图像捕获装置122-126中的一个或多个获取的多个图像的分析来识别目标车辆2310。系统100还可识别目标车辆2310行进的第二车道2340。进一步，系统100可确定第二车道2340是否不同于车辆200行进的车道2330。如果第二车道2340不同于车道2330，则系统100可构造成使车辆200能够超越目标车辆2310。

[0307] 系统100还可基于多个图像来监测目标车辆2310。例如，系统100可监测目标车辆2310的位置，诸如监测目标车辆2310相对于车辆200、相对于第一车道约束2302和/或第二车道约束2304的位置、绝对位置或相对于另一基准点的位置。另外地或替代地，监测目标车辆2310的位置可包括估算目标车辆2310的速度。

[0308] 如果系统100确定目标车辆2310正进入车辆200行进的车道2330，则系统100可致使车辆200中止对目标车辆2310的超越。例如，系统100可致使车辆200停止加速或进行制动。另外地或替代地，系统100可发布语音通知，诸如通知车辆200的驾驶员中止变道。

[0309] 图24A是与公开的实施方式相一致的存储器140的示例性框图，该存储器140可存储用于执行一个或多个操作的指令。如图24A中图示的，存储器140可存储用于识别目标车辆的一个或多个模块和本文中描述的响应。例如，存储器140可存储车道约束模块2400、目标车辆获取模块2410和动作模块2415。

[0310] 在一个实施方式中，车道约束模块2400可存储指令，当由处理单元110执行时，该指令可检测并且限定第一车道约束2302和第二车道约束2304。例如，处理单元110可执行车道偏移模块2400，以处理从至少一个图像捕获装置122-124接收的多个图像并且检测第一车道约束2302和第二车道约束2304。如上面讨论的，这可包括识别喷涂的车道线或标记物和/或测量路面的中点。处理单元110可基于上面结合图5A至图5D描述的技术来执行这种分析。

[0311] 在一些实施方式中，目标获取模块2410可存储指令，当由处理单元110执行时，该指令可检测目标车辆2310的存在并且监测目标车辆2310的位置。例如，目标获取模块2410可处理多个图像，以检测并监测目标车辆2310。另外地或替代地，当执行目标获取模块2410时，处理单元110可接收来自指示目标车辆2310的存在和/或地点的另一模块或其它系统的信息。处理单元110可基于上面结合图5A至图5D描述的技术来执行这种分析。

[0312] 处理单元110可构造成确定一个或多个对象(诸如车辆200附近的目标车辆2310)的存在。例如目标车辆2310可包括在车辆200附近行进的另一车辆，诸如汽车、卡车或摩托车。对于每个目标车辆2310，处理单元110可确定偏移轮廓。偏移轮廓可包括确定在其当前预测位置中目标车辆2310是否将位于车辆200的预定范围内。如果处理单元110确定目标车辆2310位于或将位于车辆200的预定范围内，则处理单元可确定：为使车辆200绕过目标车

IPR2025-00943
Tesla EX1002 Page 1297

辆2310，在第一车道约束810和第二车道约束820内是否有足够的空间。如果没有足够的空间，则处理单元110可执行变道。如果有足够的空间，则处理单元110可确定在车辆200与目标车辆2310之间距离缩短之前是否有足够的时间绕过目标车辆2310。如果有足够的时间，则处理单元110可启动或激活偏移机动以绕过目标车辆2310。偏移机动可由处理单元110执行，使得车辆200的移动对驾驶员而言平稳而舒适。例如，偏移摆率可大约为0.15米/秒到0.75米/秒。偏移机动可执行为使得偏移机动的最大振幅将发生在车辆200与对象825之间的间距缩短的时候。如果没有足够的时间，则处理单元110可启动制动，然后启动偏移机动。

[0313] 进一步，处理单元110可执行动作模块2415，以中止车辆200超越目标车辆2310的计划。如果目标车辆2310不在车辆200的车道2330中，则包括在动作模块1230中的指令可使车辆200能够超越目标车辆2330。然而，如果执行目标车辆获取模块2410指示目标车辆2310正变道进入车辆200的车道2330，则处理单元110可执行包括在动作模块2415中的指令以中止车辆200的变道。例如，处理单元110可向车辆200的节流系统220、制动系统230和/或转向系统240发送电子信号以中止变道。另外地或替代地，处理单元可致使语音通知发出。

[0314] 图24B图示了与公开的实施方式相一致的用于导航侵入车辆之中的车辆的示例性过程2420的流程图。过程2420可识别限定了车辆(例如，车辆200)行进的行进车道的车道约束，识别并监测目标车辆(例如，车辆2310)，并且使车辆能够超越目标车辆或者如果确定目标车辆正进入车辆的行进车道则中止超越。

[0315] 在步骤2430，图像捕获装置122、124和/或126中的至少一个可获取车辆200附近区域的多个图像。例如，包括在图像获取单元120中的相机(诸如图像捕获装置122)可捕获多个图像并且经数字连接(例如，USB、无线、蓝牙，等等)将图像发送到处理单元110。在一些实施方式中，处理单元110可经由多个数据接口接收多于一个的多个图像。例如，处理单元110可从图像捕获装置122、124、126中的每个接收多个图像，图像捕获装置122、124、126中的每个均可具有关联的数据接口用于与处理单元110通信数据。公开的实施方式并不限于任何特定的数据接口构造或协议。

[0316] 接下来，在步骤2440，处理单元110可使用经由数据接口128接收的多个图像来确定在第一车辆侧802上的第一车道约束和在第二车辆侧812上的第二车道约束。例如，作为实施步骤2440的一部分，处理单元110可执行车道约束模块910的指令。进一步，作为确定第一车道约束和第二车道约束的一部分，处理单元110可使用上面结合图5A至图5D讨论的过程中的一个或多个。

[0317] 在步骤2450，处理单元110可获取目标车辆2310。为了获取目标车辆2310，处理单元110可执行目标车辆获取模块2410。例如，处理单元110可执行指令来分析获取的图像，以识别目标车辆2310行进的第二车道2340并且确定目标车辆2310是否在车辆200行进的车道2330中行进。这种确定可基于上面结合图5A至图5D描述的技术。

[0318] 在步骤2455，处理单元110可确定目标车辆2310行进的第二车道2340是否不同于车辆200行进的车道2330。

[0319] 在步骤2460，如果第二车道2340不同于车辆200行进的车道2330，则处理单元110可使车辆200能够超越目标车辆2310。为了使车辆200能够超越目标车辆2310，处理单元可执行动作模块2415。

[0320] 接下来，在步骤2465，处理单元110可执行目标获取模块2410以监测目标车辆

IPR2025-00943
Tesla EX1002 Page 1298

2310。这可包括估算目标车辆2310的速度并且确定目标车辆2310诸如相对于车辆200的地点。

[0321] 步骤2470可包括确定目标车辆2310是否已进入车道2330。如果处理单元110确定目标车辆2310尚未进入车道2330，则在步骤2480，处理单元110可允许车辆200完成超越。

[0322] 在步骤2490，如果处理单元110确定目标车辆2310已进入车道2330(或者正进入车道2330或以其它方式处于将要把车辆2310带入车道2330中的轨迹上)，则处理单元110可执行动作模块2415以中止超越。例如，这可包括经由扬声器360向车辆200的驾驶员发布语音通知。另外地或替代地，步骤2490可包括执行包括在动作模块2415中的指令，以致使处理单元110向车辆200的节流系统220、制动系统230和/或转向系统240发送电子信号，从而中止变道。

[0323] 导航车辆以避免侵入交通

[0324] 系统100可提供驾驶员辅助功能性，用于监测车辆200附近并且响应于侵入车辆200的另一车辆的存在。监测车辆200附近可包括监测碰撞阈值。例如碰撞阈值可包括将要碰撞的时间或车辆200与其它交通之间的最小预定距离。系统100可进一步致使车辆200采取躲避动作以避免侵入交通。例如，躲避动作可包括制动、变道或以其它方式变动车辆200的路线。

[0325] 如图25中图示的，车辆200可在道路2500上行进。车辆200可包括用于检测并响应于侵入交通的系统100。例如，车辆200可确定另一车辆正进入车辆200的行进车道，或者以其它方式确定另一车辆正越过碰撞阈值。车辆200可能通过诸如变动路线、制动或加速而对侵入车辆做出反应。

[0326] 如相对于图8A至图8D讨论的并且如图25所示，车辆200可包括第一车辆侧802，第一车辆侧802距第一车道约束2502可以是第一距离805。类似地，车辆200可包括与第一车辆侧802相反的第二车辆侧812，并且第二车辆侧812距第二车道约束2504可以是第二距离815。以这种方式，第一车道约束2502和第二车道约束2504可限定车辆200可行进的车道2505。

[0327] 系统100可基于由图像捕获装置122-126中的一个或多个获取的多个图像来确定第一车道约束2502和第二车道约束2504。根据一些实施方式，第一车道约束2502和/或第二车道约束2504可由可见的车道边界(诸如标记在路面上的线)识别。另外地或替代地，第一车道约束2502和/或第二车道约束2504可包括路面的路边。根据一些实施方式，系统100可通过识别路面宽度的中点2520(或者任何其它合适的道路特征车辆200可用作导航基准)来确定第一车道约束2502和/或第二车道约束2502。当例如指定道路车道的线未被喷涂或以其它方式贴标签时，系统100可采取这种方式识别车道约束。

[0328] 检测第一车道约束2502和/或第二车道约束2504可包括处理单元110确定它们在相机坐标系统中的三维模型。例如，第一车道约束2502和/或第二车道约束2504的三维模型可由第三多项式描述。除了第一车道约束2502和/或第二车道约束2504的三维建模，处理单元110可估算运动参数，诸如车辆200的速度、偏航和俯仰速率以及加速度。可选地，处理单元可检测静态和移动车辆及其位置、航向、速度和加速度(都相对于车辆200)。处理单元110可确定道路标高模型以将从多个图像获取的信息转换为三维空间。

[0329] 车辆200的驾驶员无法假设的是，如果使交通移位到车辆200所行驶的车道2505中

IPR2025-00943
Tesla EX1002 Page 1299

是安全的，则在相邻或附近相邻车道中行进的其它交通将仅变道。由此，可期望考虑并监测在近旁车道中行进的交通，以检测对于车辆200的车道2505的侧向侵入，使得车辆200可对侧向侵入做出反应以避免碰撞。系统100可基于多个图像来识别侵入车辆2510并且确定侵入车辆2510正接近车辆200。然后，系统100可产生动作来考虑侵入车辆。例如，系统100可致使车辆200维持当前速度并且在第一车道约束2502和第二车道约束2504内行进，使得在侵入车辆2510接近的第一车辆侧802上的第一距离2530大于第二距离2540。

[0330] 处理单元110可构造成确定车辆200附近的一个或多个侵入车辆2510的存在。对于每个侵入车辆2510，处理单元110可确定偏移轮廓。偏移轮廓可包括确定在其当前预测位置中侵入车辆2510是否将位于车辆200的预定范围内。如果处理单元110确定侵入车辆2510位于或将位于车辆200的预定范围内，则处理单元可确定：为使车辆200绕过侵入车辆2510，在第一车道约束2502和第二车道约束2504内是否有足够的空间。如果没有足够的空间，则处理单元110可执行变道。如果有足够的空间，则处理单元110可确定在车辆200与侵入车辆2510之间距离缩短之前是否有足够的时间绕过侵入车辆2510。

[0331] 例如，如果有足够的时间，则处理单元110可启动或激活偏移机动以绕过侵入车辆2510。偏移机动可由处理单元110执行，使得车辆200的移动对驾驶员而言平稳而舒适。例如，偏移摆率可大约为0.15米/秒到0.75米/秒。偏移机动可执行为使得偏移机动的最大振幅将发生在车辆200与侵入车辆2510之间的间距缩短的时候。如果没有足够的时间，则处理单元110可启动制动(例如，经由向制动系统230发送电子信号)，然后启动偏移机动(例如，通过向节流系统220和/或转向系统240发送电子信号)。

[0332] 系统100可基于由图像捕获装置122-126中的一个或多个获取的多个图像来确定：侵入车辆2510已越过至少一个碰撞阈值。例如，至少一个碰撞阈值可包括车辆200与侵入车辆2510之间的最小预定距离2550。例如，如图25中图示的，侵入车辆2510示出为超过最小预定距离2550。另外地或替代地，至少一个碰撞阈值可包括车辆200与侵入车辆2510将要碰撞的时间。基于时间的碰撞阈值可基于下列来确定：车辆200与侵入车辆2510之间的距离、车辆200和/或侵入车辆2510的速度、车辆200和/或侵入车辆2510的侧向速度和/或车辆200和/或侵入车辆2510的加速度。基于碰撞时间的碰撞阈值可设定为允许系统100有足够的时间来有效进行躲避动作以避免碰撞。

[0333] 如果系统100确定至少一个碰撞阈值已被越过，则系统100可致使车辆200采取躲避动作。例如，这可包括致使车辆200的其它子系统操作。由此，系统100可操作成通过与节流系统220、制动系统230和/或转向系统240通信来致使车辆200的速度增加，以改变车辆200的速度和/或方向。

[0334] 躲避动作可包括致使车辆200变道。可期望在例如车辆200在多车道高速公路上行进的情况下响应，并且系统100确定：变道远离侵入车辆2510将避免碰撞。躲避动作可包括使车辆200变动其路线。另外地或替代地，躲避动作可包括变动车辆200的路线。例如，为了采取躲避动作，处理单元110可向节流系统220、制动系统230和/或转向系统240发送电子信号。

[0335] 系统100还可构造成在躲避动作已完成之后致使车辆200采取动作。例如，系统100可致使车辆200返回到车辆200行进的车道2505的中点2520。另外地或替代地，系统100可在响应于侵入车辆2510制动之后致使车辆200恢复预设速度。

IPR2025-00943
Tesla EX1002 Page 1300

[0336] 图26是与公开的实施方式相一致的存储器140的示例性框图，该存储器140可存储用于检测并且响应于车辆的交通侧向侵入的指令。如图26中图示的，存储器140可存储车道约束模块2610、侵入车辆获取模块2620和动作模块2630。

[0337] 在一个实施方式中，车道约束模块2610可存储指令，当由处理单元110执行时，该指令可检测并且限定第一车道约束2502和第二车道约束2504。例如，处理单元110可执行车道偏移约束模块2610，以处理从至少一个图像捕获装置122-126接收的多个图像并且检测第一车道约束2502和第二车道约束2504。如上面讨论的，这可包括识别喷涂的车道线和/或测量路面的中点。处理单元110可基于上面结合图5A至图5D描述的技术来执行这种分析。

[0338] 在一个实施方式中，侵入车辆获取模块2620可存储指令，当由处理单元110执行时，该指令可检测侵入车辆2510的存在并且监测侵入车辆2510。例如，侵入车辆获取模块2620可包括用于确定侵入车辆2510关于车辆200的相对地点的指令。在一些实施方式中，这可包括监测侵入车辆200的速度。根据一些实施方式，侵入车辆获取模块2620可确定侵入车辆2510是否已超过碰撞阈值。例如，碰撞阈值可以是车辆200与侵入车辆2510之间的最小预定距离。如果侵入车辆2510比最小预定距离更接近车辆200，则处理单元110可确定侵入车辆2510的确在侵入并且执行指令以避免与侵入车辆2510碰撞。

[0339] 另外地或替代地，碰撞阈值可包括碰撞时间。基于时间的碰撞阈值可基于下列来确定：车辆200与侵入车辆2510之间的距离、车辆200和/或侵入车辆2510的速度、车辆200和/或侵入车辆2510的侧向速度和/或车辆200和/或侵入车辆2510的加速度。基于碰撞时间的碰撞阈值可设定为允许处理单元110和系统100有足够的时间来有效进行躲避动作以避免碰撞。

[0340] 进一步，处理单元110可执行动作模块2630，以响应于侵入车辆2510的检测和/或越过碰撞阈值。例如，动作模块2630可包括这样的指令，该指令用于致使车辆200维持当前速度并且在第一车道约束2502和第二车道约束2504内行进，使得第一距离2530大于第二距离2540。第一距离2530可位于车辆200的与侵入车辆2510相同的那一侧上，使得车辆200在仍位于第一车道约束2502和第二车道约束2504内的同时更远离侵入车辆2510。如果处理单元110确定碰撞阈值已被越过，则处理单元110可执行动作模块2630，以致使车辆200采取躲避动作，诸如变动路线、制动和/或变道。动作模块2630还可包括这样的指令，在完成躲避动作之后，该指令致使车辆200恢复路线。例如，在制动或返回到车道2505的中央之后，处理单元110可致使车辆200恢复预设速度。

[0341] 图27图示了与公开的实施方式相一致的检测并且响应于车辆的交通侧向侵入的示例性过程2720的流程图。过程2720可识别限定车辆行进的行进车道的车道约束，确定侵入车辆2510是否接近，并且致使车辆200维持当前速度并在第一车道约束2501和第二车道约束2504内行进，使得在车辆200的侵入车辆2510接近的那一侧上的第一距离2530大于第二距离2540。

[0342] 在步骤2730，至少一个图像捕获装置122、124和/或126可获取车辆200附近区域的多个图像。例如，包括在图像获取单元120中的相机(诸如图像捕获装置122)可捕获多个图像并且经数字连接(例如，USB、无线、蓝牙，等等)将图像发送到处理单元110。在一些实施方式中，处理单元110可经由多个数据接口接收多于一个的多个图像。例如，处理单元110可从图像捕获装置122、124、126中的每个接收多个图像，图像捕获装置122、124、126中的每个均

IPR2025-00943
Tesla EX1002 Page 1301

可具有关联的数据接口用于与处理单元110通信数据。公开的实施方式并不限于任何特定的数据接口构造或协议。

[0343] 接下来，在步骤2740，处理单元110可使用经由数据接口128接收的多个图像来确定在第一车辆侧802上的第一车道约束和在第二车辆侧812上的第二车道约束。例如，作为实施步骤2740的一部分，处理单元110可执行车道约束模块2610的指令。进一步，作为确定第一车道约束和第二车道约束的一部分，处理单元110可使用上面结合图5A至图5D讨论的过程中的一个或多个。

[0344] 在步骤2750，处理单元110可确定侵入车辆2510是否接近。为了做出这种确定，处理单元110可执行侵入车辆获取模块2620。侵入车辆获取模块2620可包括处理图像以检测其它交通的指令。这可包括在车辆200的特定范围内识别车辆，诸如邻近车辆200的那些车辆。这些邻近车辆可能是侵入车辆2510。

[0345] 一旦侵入车辆2510被识别，处理单元就可确定关于车辆2510的更多信息。例如，在步骤2555，处理单元110可确定目标车辆2330是否已超过碰撞阈值。例如，碰撞阈值可以是车辆200与侵入车辆2510之间的最小预定距离。如果侵入车辆2510比最小预定距离更接近车辆200，则处理单元110可确定侵入车辆2510的确在侵入并且执行指令以避免与侵入车辆2510碰撞。

[0346] 另外地或替代地，碰撞阈值可以是碰撞时间。基于时间的碰撞阈值可基于下列来确定：车辆200与侵入车辆2510之间的距离、车辆200和/或侵入车辆2510的速度、车辆200和/或侵入车辆2510的侧向速度和/或车辆200和/或侵入车辆2510的加速度。基于碰撞时间的碰撞阈值可设定为允许处理单元110和系统100有足够的时间来有效进行躲避动作以避免碰撞。

[0347] 在步骤2760，处理单元110可致使车辆200维持当前速度并且在第一车道约束2502和第二车道约束2504内行进。为了致使车辆200以这种方式操作，处理单元110可执行动作模块2630。另外地或替代地，步骤2760可包括执行动作模块1230上的指令，以致使处理单元110向车辆200的加速器2610、制动器2620和/或转向系统发送电子信号，致使车辆200维持当前速度并且在第一车道约束810和第二车道约束820内行进。

[0348] 在步骤2755，如果处理单元110确定碰撞阈值已被超过，则处理单元110可采取躲避动作。躲避动作可包括致使车辆200变道。可期望响应的是，例如，车辆200在多车道高速公路上行进，并且系统100确定变道远离侵入车辆2510将避免碰撞。躲避动作可包括使车辆200变动其路线。另外地或替代地，躲避动作可包括变动车辆200的路线。例如，处理单元110可通过向车辆200的节流系统220、制动系统230和/或转向系统240发送电子信号来采取躲避动作。在完成躲避动作之后，在步骤2780，处理单元110可控制车辆200恢复其原先的导航(例如，恢复导航至期望的地点)。

[0349] 车辆导航的多阈值反应区

[0350] 系统100可提供驾驶员辅助功能性，用于产生车辆200的响应，诸如制动、加速、切换车道、转弯和/或其它导航响应。例如，如果系统100检测存在着位于车辆200之前以慢于车辆200的速度行进的另一车辆，系统100可通过制动致使车辆200以减小其速度。系统100可施加制动轮廓，该制动轮廓考虑到风险(例如，碰撞的风险)的即时性，为驾驶员提供自然驾驶感觉。

IPR2025-00943
Tesla EX1002 Page 1302

[0351] 例如，取决于与另一车辆碰撞的情形和确定性风险，系统100可产生各种不同响应。如果目标车辆在前面足够远并且碰撞的确定性风险较低，则系统100响应于前面目标车辆的速度或方向改变可不采取动作。当目标车辆可能更近时或者当碰撞的风险上升超过预定水平时，系统100可产生响应，诸如制动，等等。当碰撞风险甚至更高时（例如，系统100确定碰撞将迫在眉睫而无躲避动作），系统100可致使车辆100采取躲避动作，诸如最大或近乎最大制动、方向改变，等等。

[0352] 图28A和图28B是与公开的实施方式相一致的主车辆200a在道路2800上的概略表示。在一些实例中，道路2800可分为两个车道，车道2810和车道2820。诸如目标车辆200b的目标对象可在同一车道（例如，如图28A和图28B所示的车道2820）的主车辆200a之前行进或在不同车道中。如图28A描绘的示例所示，主车辆200a和目标车辆200b可分离距离d_1。在图28A描绘场景之后的给定时间里，如图28B所示，主车辆200a和目标车辆200b可分离距离d_2，其中d_2小于d_1。例如，主车辆200a能以大于目标车辆200b的速度行进，致使车辆之间的距离从d_1降至d_2。

[0353] 作为另一示例，主车辆200a和目标车辆200b能以相同的速度行进，直到目标车辆200b通过制动减少其速度，或者主车辆200a通过加速增加其速度，致使车辆之间距离从d_1降至d_2。作为另一示例，目标车辆200b可从主车辆200a的一侧侵入。在一些情形下，如果目标车辆200b在主车辆200a之前或者远离主车辆200a足够距离（例如，五米以上，10米以上，等等）并且移向主车辆200a，系统100可不采取任何动作。然而，如果目标车辆200b继续移动而更接近主车辆200a，则系统100可致使主车辆200a制动、改变速度和/或变道。例如，如果目标车辆200b在预定阈值（例如，五米内）内，则主车辆200a可采取动作。涉及主车辆200a和目标车辆200b的相对速度和位置可以有不同的场景，并且公开的实施方式并不限于任何特定情景。

[0354] 在一些实施方式中，多阈值反应区可向用户提供自然感觉驾驶经历。例如，如上面讨论的，如果碰撞风险较低，则系统100可忽略目标车辆200b的移动。结果是，当不存在碰撞风险时（例如，当目标车辆200b在100米之前时，主车辆200a将不需要模仿目标车辆200b的每个运动），主车辆200a可不响应于目标车辆200b的移动或动作。然而，如果目标车辆200b接近主车辆200a和/或如果系统100确定碰撞风险较高，则可采取响应式机动，通过增加主车辆200a与目标车辆200b之间的空间使得用户感觉更安全。如果系统100确定撞击可能迫在眉睫，则系统100可采取更激进的动作而不会致使用户感到惊讶，因为为避免碰撞该动作可能是必要的。

[0355] 因此，在一些实施方式中，系统100可将不同的反应区与不同程度的响应动作关联。例如，在一些实施方式中，当目标车辆200b在第一阈值内时（例如，当目标车辆200b远离主车辆200a时），系统100可不采取动作；当目标车辆200b在第二阈值内时（例如，当目标车辆200b更接近主车辆200a时），系统100可采取中等动作；并且当目标车辆200在第三阈值内时（例如，当目标车辆200b足够接近主车辆200a存在碰撞风险时），系统100可采取更激进的动作。在与第一阈值关联的第一反应区中，当与目标车辆200b关联的距目标的估算距离比预定安全距离大两倍时，系统100可不采取动作。在一些实施方式中，在与目标车辆200b关联的距目标的估算距离落入预定安全距离的1.5倍到2.5倍的范围内的情况下，可应用第一反应区。

IPR2025-00943
Tesla EX1002 Page 1303

[0356] 距目标的估算距离可以是主车辆200a与目标车辆200b之间的当前距离加上该距离经预测时间(例如,在0.1秒到1.0秒的范围内,取决于与系统100关联的因素,诸如速度跟踪控制环、致动器和车辆动态)后的预期改变。预期距离改变可计算为:$(v_{目标}-v_{主})*t_{预测}+((a_{目标}-a_{主})*t_{预测}^{1/2})$,其中$v_{目标}$和$v_{主}$分别对应于目标车辆200b和主车辆200a的速度,$a_{目标}$和$a_{主}$分别对应于目标车辆200b和主车辆200a的加速度,并且$t_{预测}$对应于预测时间。选择逼近系统100的整体反应时间的预测时间可考虑以下事实,即主车辆200a的期望导航响应可不立即实施,从而可为主车辆200a的驾驶员和/或乘客提供平稳的自然驾驶经历。

[0357] 预定安全距离可定义为:最大{$d_{静态}$,$d_{动态}$}。可落入2米到10米范围内的$d_{静态}$可表示车辆停止和/或低速移动时主车辆200a与另一车辆(诸如目标车辆200b)之间的期望安全距离。$d_{动态}$可表示车辆以大于低速的速度移动时的期望安全距离,并且可计算为:$t_{安全}$*最小($v_{目标}$,$v_{主}$),其中$v_{目标}$和$v_{主}$分别对应于目标车辆200b和主车辆200a的速度。$t_{安全}$可落入0.5秒到2.5秒的范围内,并且可根据驾驶员的喜好由主车辆200a的驾驶员经由用户接口170调整。由此,驾驶员能够控制主车辆200a与其它车辆之间系统100维持的距离。

[0358] 在与第二阈值关联的第二反应区中,系统100可将目标车辆200b放置在安全区中并且向由目标车辆200b执行的不同机动施加不同的权重。例如,在目标车辆200b可接近安全区的下边界时,例如,与目标车辆200b关联的距目标的估算距离可接近预定安全距离(上面结合第一反应区讨论过),系统100可以向机动减速分配更大的权重(例如,模仿更多)并且向机动加速分配更小的权重(例如,模仿更少)。在一些实施方式中,下边界可落入预定安全距离的0.5倍到1.5倍的范围内(取决于驾驶员的喜好)。反过来,在目标车辆200b接近安全区的上边界时,例如,与目标车辆200b关联的距目标的估算距离可接近预定安全距离的两倍,系统100可以向机动减速分配更小的权重(例如,模仿更少)并且向机动加速分配更大的权重(例如,模仿更多)。在一些实施方式中,下边界可落入预定安全距离的1.5倍到2.5倍的范围内(取决于驾驶员的喜好)。

[0359] 在与第三阈值关联的第三反应区中,当与目标车辆200b关联的距目标的估算距离小于预定安全距离时,系统100可将目标车辆200b放置在危险区中。在该情景下,系统100可致使主车辆200a模仿(例如,立即模仿)目标车辆200b执行的任何减速机动(同时不考虑任何加速机动)。在一些实施方式中,在与目标车辆200b关联的距目标的估算距离落入预定安全距离的0.5倍到1.5倍的范围内(取决于驾驶员的喜好)的情况下,可应用第一反应区。

[0360] 在一些实施方式中,系统100可产生主车辆200a的响应,诸如制动。如图28B中指示的,例如,主车辆200a与目标车辆200b之间的距离可从d_1降至d_2。取决于距离d_2并且延伸至在碰撞之前的时间,系统100可致使主车辆200a减少其速度以避免与目标车辆200b碰撞。另外,还基于例如在碰撞之前的时间,当致使主车辆200a减少其速度时,系统100可施加一个或多个制动轮廓。由此,如果在碰撞之前的时间(和/或距离d_2)超过第一预定阈值,则系统100可致使主车辆200a以渐进的方式制动、变道和/或改变速度,而如果在碰撞之前的时间(和/或距离d_2)低于第二预定阈值,则系统100可致使主车辆200a制动、变道和/或以更快速的方式改变速度。

[0361] 图29是与公开的实施方式相一致的存储器的示例性框图,该存储器被构造成存储用于执行一个或多个操作的指令。如图29所示,存储器140可存储目标监测模块2902、拦截时间模块2904和动作响应模块2906。公开的实施方式并不限于任何特定构造的存储器140。

IPR2025-00943
Tesla EX1002 Page 1304

进一步，处理单元110可执行存储在任何模块2902-2906(包括在存储器140中)中的指令。

[0362] 在一个实施方式中，目标监测模块2902可存储软件指令(诸如计算机视觉软件)，当由处理单元110执行时，该软件指令识别并且监测出现在表示主车辆200附近环境的多个图像中的目标对象(诸如目标车辆200b)。出于本公开的目的，目标对象也可被称为目标车辆200b，但是目标对象可以是行人、道路危险(例如，大的碎片)，等。多个图像可由图像捕获装置122、124和126中的一个或多个获取。处理单元110可执行目标监测模块2902，以使用上面结合图5A至图5D描述的技术通过例如分析多个图像来识别并且监测目标车辆200b。由此，目标监测模块2902可跟踪不同的变量，诸如与目标车辆200b关联的运动(例如，行进方向和轨迹)、速度和加速度，以及主车辆200a与目标车辆200b之间的距离。这样的信息可不断地或者以预定的间隔跟踪。

[0363] 在一个实施方式中，拦截时间模块2904可存储软件指令，当由处理单元110执行时，该软件指令确定主车辆200a和目标车辆200b之间的拦截时间的指示器。拦截时间可指在主车辆200a与目标车辆200b接触(例如，碰撞)之前的时间量。处理单元110可执行拦截时间模块2904，以基于分析出图像捕获装置122、124和126中的一个或多个获取的图像来确定拦截时间的指示器。例如，拦截时间模块2904可使用例如目标车辆200b(相对于主车辆200a)的位置、速度和/或加速度的一系列基于时间的观察(诸如卡尔曼滤波器或线性二次估算(LQE))来做出确定。卡尔曼滤波器可基于目标车辆200b的尺度的测量结果，其中尺度测量结果与拦截时间成比例。

[0364] 在一个实施方式中，动作响应模块2906可存储软件指令(诸如计算机视觉软件)，当由处理单元110执行时，该软件指令在主车辆200a中产生一个或多个响应。所述响应可基于目标车辆200b的识别和监测(例如，经由执行目标监测模块2902来执行)和/或主车辆200a与目标车辆200b之间的拦截时间的指示器(例如，经由执行拦截时间模块2904确定)。处理单元110可通过例如(例如，经由CAN总线)向主车辆200a的节流系统220、制动系统230和/或转向系统240发送电子信号来在主车辆200a中产生一个或多个响应，以制动、加速或触发主车辆200a的转弯、车道移位和/或方向改变。处理单元110可使用扩展的卡尔曼滤波器来估算与目标车辆200b和主车辆200a关联的数据，诸如：速度、加速度、转弯率、俯仰角、俯仰速率、目标车辆200b相对于主车辆200a的位置和航向、目标车辆200b的大小、道路几何形状，等。例如，处理单元110可接收来自主车辆200a上的部件(诸如里程表、惯性测量单元(IMU)、图像处理器190，等)的测量结果。由此，处理单元110可使用从运动的运动学公式得到的状态公式、刚性约束和运动假设来计算与目标车辆200b和/或主车辆200a关联的数据，其中公式的输入可包括来自主车辆200a上的部件的测量结果。处理单元110还可使用上面结合图4至图7描述的技术，以产生主车辆200a的一个或多个响应。

[0365] 图30是示出与公开的实施方式相一致的用于产生主车辆的响应的示例性过程3000的流程图。在步骤3010，处理单元110可经由位于处理单元110与图像获取单元120之间的数据接口128接收多个图像。例如，包括在图像获取单元120中的一个或多个相机可捕获主车辆200a前方(例如，或车辆的侧方或后面)的区域的多个图像并且经数据连接(例如，数字、有线、USB、无线、蓝牙，等等)将图像发送到处理单元110。

[0366] 在步骤3020，处理单元110可执行目标监测模块2902，以识别多个图像内的目标对象(诸如目标车辆200b)。例如，目标监测模块2902可使用上面结合图5A和图5B描述的技术

IPR2025-00943
Tesla EX1002 Page 1305

来检测图像中的车辆的存在。基于图像数据、位置数据(例如,GPS地点信息)、地图数据、主车辆200a的偏航率、车道或其它道路标记、速度数据和/或包括在主车辆200a中的传感器的数据,多个车辆(和/或对象)可在与主车辆200a相同的车道或不同的车道中被检测,并且距主车辆200a的距离是变化的。在检测多个车辆的情景下,目标监测模块2902可将目标车辆200b确定为在与主车辆200a相同的车道中行进和/或最接近主车辆200a的车辆。

[0367] 在步骤3030,处理单元110可执行目标监测模块2902,以监测在步骤3020识别的目标车辆200b。例如,目标监测模块2902可跟踪与目标车辆200b关联的信息,诸如运动(例如,行进方向和轨迹)、速度和加速度,以及主车辆200a与目标车辆200b之间的距离。目标监测模块2902可使用上面结合图5A至图5D描述的技术通过分析多个图像来不断地或者以预定的间隔跟踪这样的信息。在步骤3040,处理单元110可执行拦截时间模块2904,以确定主车辆200a与目标车辆200b之间的拦截时间的指示器。拦截时间模块2904可通过跟踪在一个或多个时间周期上目标车辆200b(相对于主车辆200a)的位置、速度和/或加速度来做出确定。拦截时间模块2904可使用上面结合图5A和图5B描述的技术来跟踪这样的信息。

[0368] 在步骤3050,处理单元110可执行动作响应模块2906,以产生主车辆200a的一个或多个响应。例如,处理单元110可(例如,经由CAN总线)向主车辆200a的节流系统220、制动系统230和/或转向系统240发送电子信号,以制动、加速或触发主车辆200a的转弯、车道移位和/或方向改变。进一步,一个或多个致动器可控制节流系统220、制动系统230和/或转向系统240。例如,处理单元110可发送电子信号,致使系统100物理地踩下制动器达预定的量或者部分地松开车辆200的加速器。进一步,处理单元110可发送致使系统100车辆200沿特定方向转向的电子信号。这样的响应可基于(在步骤3030执行的)目标车辆200b的监测和/或(在步骤3040确定的)主车辆200a与目标车辆200b之间的拦截时间的指示器。在一些实施方式中,拦截时间可与一个或多个预定阈值(例如,第一和第二拦截阈值)比较。例如,如果拦截时间的指示器指示目标对象在特定拦截阈值内(例如,在第二拦截阈值内),则所述响应可包括紧急避让动作。紧急避让动作可包括变道和/或紧急制动。

[0369] 阈值可表达为时间的单位。由此,例如,如果拦截时间的指示器指示目标车辆200b在第一阈值之外(例如,在主车辆200a与目标车辆200b之间的碰撞之前的时间超过第一阈值),则动作响应模块2906可基于与目标车辆200b关联的平均运动轮廓来在主车辆200a中产生响应。与目标车辆200b关联的平均运动轮廓可表示目标车辆200b针对特定阈值的平均位置和/或平均速度。在一些实施方式中,如果目标车辆200b为某一距离和/或以某一速度行进,则系统100可不采取动作。例如,如果目标车辆200b在主车辆200a之前或者距离足够远离主车辆200a(例如,五米、10米,等等),则系统100可不采取任何动作,除非目标车辆200b更接近主车辆200。

[0370] 如果拦截时间的指示器指示目标车辆200b在第一阈值与第二阈值之间(例如,在主车辆200a与目标车辆200b之间的碰撞之前的时间落入第一阈值与第二阈值之间),则动作响应模块2906可基于与目标车辆200b关联的实际运动来产生主车辆200a的响应。实际运动可表示目标车辆200b的位置和/或速度(例如,并未平均或延迟),如果拦截时间的指示器指示目标车辆200b在第二阈值内(例如,在主车辆200a与目标车辆200b之间的碰撞之前的时间低于第二阈值),则动作响应模块2906可产生主车辆200a的紧急动作(诸如变道或紧急制动)。

IPR2025-00943
Tesla EX1002 Page 1306

[0371] 图31是示出与公开的实施方式相一致的用于拒绝在主车辆中产生响应的示例性过程3100的流程图。在步骤3110，如上面结合图30的步骤3010描述的，处理单元110可经由位于处理单元110与图像获取单元120之间的数据接口128接收多个图像。

[0372] 在步骤3120，如上面结合图30的步骤3020描述的，处理单元110可执行目标监测模块2902，以通过分析图像将目标对象(例如，目标车辆200b)定位在步骤3010接收的图像内。在步骤3125，处理单元110可执行目标监测模块2902，以使用上面结合图30的步骤3030描述的技术来确定目标对象(例如，目标车辆200b)是否在不同于主车辆200a的车道中行进。如果确定目标车辆200b在不同于主车辆200a的车道中行进(步骤3125，是)，则在步骤3130，处理单元110可拒绝在主车辆200a中产生响应。替代地，如果系统100确定目标车辆200b不构成迫在眉睫的碰撞风险(例如，目标车辆200b在与主车辆200a非相交的轨线上或者潜在的将要碰撞的时间超过阈值，使得目标对象在目前的时间里不构成迫在眉睫的风险)，在步骤3130，处理单元110可拒绝在主车辆200a中产生响应。否则(步骤3125，否)，过程3100可结束且不发生响应。

[0373] 使用多区段制动轮廓来停止自主驱动车辆

[0374] 系统100可提供驾驶员辅助功能性，用于监测并且响应于环境条件以控制车辆200的制动。例如，系统100可基于多个图像来检测触发车辆200停止的对象。例如，该对象可包括车辆200附近的另一车辆、交通灯和/或交通标志。系统100可使用制动轮廓来减速和/或停止车辆200。根据一些实施方式，制动轮廓可包括在不同阶段执行减速的多个区段。例如，制动轮廓可致使系统100以降速的减速度停止车辆200。

[0375] 通过提供包括多个区段的制动轮廓，系统100可致使车辆200以对驾驶员和/或乘客而言感觉自然的方式操作。例如，当驾驶员在交通灯处停止车辆时，从施加的时刻直到车辆停止的时刻，制动器通常不以恒速来施加。以这样的方式操作制动器导致终止时突然停止。相反，驾驶员可首先平缓地施加制动器，然后增加压力，进而在车辆到达其目标停止地点时，驾驶员可降低制动水平，直到在汽车停止的时刻到达零(或零附近)。以这种方式操作车辆可提供自然停止过程，即，驾驶员可首先致使逐步减速，所以车辆不会经历突然初始减速然后终止，使得车辆不会突然停止。

[0376] 系统100可提供可根据对车辆200的驾驶员和/或乘客而言感觉自然的方法来实施制动的自主车辆导航。例如，在一些实施方式中，系统100可致使图像捕获装置122-126中的一个或多个获取车辆200附近区域的多个图像。处理装置110可经由数据接口128接收多个图像，并且基于对多个图像的分析来识别用于停止车辆200的触发器(例如，对象、交通灯和/或交通标志)。基于被识别的触发器，系统100可根据制动轮廓致使车辆200停止。在一些实施方式中，如基于对多个图像的分析而确定的，制动轮廓可包括多个区段(例如，包括：与第一减速率关联的第一区段；包括小于第一减速率的第二减速率的第二区段；以及在接近目标停止地点时使制动水平降低的第三区段)。

[0377] 图32A图示了在包括对象3220的道路3210上行进的车辆200。系统100可识别触发器(例如，对象3220)以使车辆200相对于触发器的检测而停止。例如，在触发器被检测到之后系统100可首先使车辆200减速，并且在到达对象3220之前使车辆200停止。

[0378] 处理单元110可构造成基于由图像捕获装置122-126中的一个或多个获取的多个图像来确定用于停止车辆200的触发器。处理单元110可经由数据接口128从图像捕获装置

IPR2025-00943
Tesla EX1002 Page 1307

122-126接收图像。根据一些实施方式，触发器可包括对象3220。例如，对象3220可以是道路2810上的另一车辆。根据一些实施方式，如果对象3220在车辆200的某一附近区域内，则处理单元110可识别触发器。例如，识别触发器可取决于对象3220与车辆200之间的相对距离。进一步，这种识别可选地考虑对象3220与车辆200之间的相对速度和/或相对加速度。

[0379] 另外地或替代地，触发器可包括交通标志，诸如图32A中示出的交通标志3240。例如，交通标志3240可以是停止标志和/或限速标志。根据一些实施方式，处理单元110可处理多个图像以确定交通标志3240的内容，作为确定触发器是否存在的一部分。例如，如果车辆200超过限速，则处理单元110可基于由交通标志3240规定的限速来确定触发器的存在。作为另一示例，如果黄灯亮和/或红灯亮，则处理单元110可基于交通灯2830来确定触发器的存在。

[0380] 基于多个图像，处理单元110可确定目标停止地点3250，一旦执行制动轮廓，车辆200就将停止在目标停止地点3250处。到达目标停止地点3250可伴随着车辆200的制动压力的释放。目标停止地点3250可基于交通灯3230和/或交通标志3240的地点，使得车辆200符合交通灯3230和/或交通标志3240。替代地，目标停止地点3250可基于对象3220，使得车辆200在到达对象3220之前停止。因此，目标停止地点3250可基于对象3220的地点，以避免车辆200与对象3220之间的碰撞。在一些实施方式中，目标停止地点3250的更新可基于从多个图像获得的新的信息，诸如新识别的目标、新的可见障碍、新的可见标志，等等。

[0381] 图32B图示了与公开的实施方式相一致的示例性制动轮廓。在图32B示出的示例中，制动轮廓包括四个区段。然而，区段的数目是示例性的，并且与公开的实施方式相一致的制动轮廓可包括任何适当数目的区段(例如，2个、3个、4个、5个、6个等等区段)。

[0382] 在区段1中，系统100可准备可能的制动情景。例如，在该区中，系统100可接收关于即将到来的路口(例如，具有交通灯的路口)的信息(例如，地图数据和/或图像数据)。系统100可逐渐减小车辆200的速度(例如，取决于路口的特性，车辆200的速度可逐渐减小至约70-50km/hr)。系统100可产生速度调整，例如，大约在枢纽之前的100-200米平均减速度约0.2-0.4米/秒^2(最大约0.5米/秒^2)。通过在区段1中降低车辆200的速度，系统100可为车辆200的驾驶员和/或乘客提供系统100知晓接近路口的置信度。

[0383] 在区段2中，系统100可施加强大的制动。在该区段中，系统100可基于对图像数据的分析来接收信息，例如，交通灯变红和/或在车辆200行进线中的路口处存在停止车辆。系统100可致使车辆200在该区中经历显著的速度减小，同时距停止线或停止车辆的距离仍较大。这样做可为车辆200的驾驶员和/或乘客提供舒适的感觉，系统100将具有足够多的空间来完成制动机动。该区可远离停止线大约30-100米，并且平均减速度可约为1.5-2.5米/秒^2(最大约为3.5米/秒^2)。

[0384] 在区段3中，系统100可做出适度制动调整。例如，在该区段中，系统100可根据距停止线或停止车辆的剩余距离并且基于车辆200的当前速度和减速度来调整车辆200的速度。通过做出这种调整，系统100可为车辆200的驾驶员和/或乘客提供如下指示：系统100正释放大部分制动功率，同时速度仍减小，但慢于之前。该区可远离停止线或停止车辆大约5-30米，并且平均减速度可约为0.5-1.5米/秒^2(最大约为2米/秒^2)。

[0385] 在区段4中，系统100可做出小幅制动调整。在该区段中，系统100能以逐渐变为零的非常低的速度缩短距停止线或停止车辆的剩余距离。在该区中，驾驶员和/或乘客可感觉

IPR2025-00943
Tesla EX1002 Page 1308

车辆200缓缓地滑动到位(例如,在停止线处或在另一车辆后面)。该区可构成距停止线或停止车辆的大约最后5-7米,并且平均减速度可约为0.3-0.5米/秒^2(最大约为1米/秒^2)。

[0386] 图33是与公开的实施方式相一致的存储器140和/或150的示例性框图,该存储器140和/或150可存储用于执行一个或多个操作的指令。如图33中图示的,存储器140可存储用于执行本文中描述的触发器检测和响应的一个或多个模块。例如,存储器140可存储触发器识别模块3300和动作响应模块3310。

[0387] 触发器识别模块3300可存储指令,当由处理单元110执行时,该指令可检测用于停止车辆200的触发器的存在。例如,触发器识别模块3300可处理从至少一个图像捕获装置122-124接收的多个图像,以检测触发器的存在。如上面讨论的,这可包括识别对象3220、交通灯3230和/或交通标志3240。

[0388] 动作响应模块3310可存储指令,当由处理单元110执行时,该指令可响应于触发器的存在。例如,动作响应模块3310可执行控制,以减速和/或停止车辆200。例如,动作响应模块3310可执行制动轮廓,以减速和/或停止车辆200。动作响应模块3310可基于制动轮廓的数目并且基于例如环境因素、被检测触发器的类型和车辆200的速度来确定执行多个制动轮廓中的一个。为了基于被检测触发器将车辆200导航至停止,处理单元110可向车辆200的节流系统220、制动系统230和/或转向系统240中的一个或多个发送电子信号。进一步,一个或多个致动器可控制节流系统220、制动系统230和/或转向系统240。例如,处理单元110可发送电子信号,致使系统100物理地踩下制动器达预定的量。

[0389] 图34图示了与公开的实施方式相一致的用于导航车辆200的过程3400。过程3000可识别用于停止车辆200的触发器(例如,对象、车辆、交通灯、交通标志,等等)并且响应于触发器的检测来执行制动轮廓。过程3400可检测各种类型的触发器,以确定执行多个制动轮廓中的哪一个。

[0390] 制动轮廓可基于被检测触发器的类型和/或与被检测触发器相关的其它环境条件(例如,距对象的距离)。例如,在一些实施方式中,制动轮廓可包括多个区段(例如,三个区段)。在三区段制动轮廓中,制动轮廓的第一区段可与第一减速率关联,制动轮廓的第二区段可包括小于第一减速率的第二减速率,并且制动轮廓的第三区段可包括在接近目标停止地点时致使制动水平降低,这基于对由图像捕获装置122-126中的一个或多个获取的多个图像的分析来确定。例如,在一些实施方式中,第一区段可致使车辆200逐步增加减速度,第二区段可致使车辆200恒速减速,并且第三区段可致使车辆200将其减速率逐步降低至零(或附近零)。

[0391] 在步骤3410,过程3400可使用图像捕获装置122-126中的至少一个来获取车辆200附近区域的多个图像。例如,处理单元110可经由数据接口128接收多个图像。处理单元110可构造成基于获取的图像来确定存在一个或多个对象和/或位于道路3210之上或附近的标志。例如,处理单元110可构造成确定对象3220、交通灯3230和/或交通标志3240的存在。处理单元110可进一步构造成读取交通标志3240的内容。例如,处理单元110可构造成在限速标志与停止标志之间进行区分。进一步,处理单元110可构造成基于交通标志3240来处理图像以确定限速。

[0392] 在步骤3420,过程300可包括用于停止车辆200的触发器。例如,处理单元110可基于对象3220在车辆200附近区域内的存在来识别触发器。另外地或替代地,处理单元110可

IPR2025-00943
Tesla EX1002 Page 1309

基于交通灯3230的当前信号状态来识别触发器。根据一些实施方式，如果不是绿灯，则处理单元110可确定交通灯3230构成触发器。根据一些实施方式，例如，如果交通标志3240是停止标志，则处理单元110可确定交通标志3240构成触发器。在一些实施方式中，如果交通标志3240是限速标志并且车辆200超过限速，则交通标志3240可构成触发器。进一步在一些实施方式中，处理单元110可致使系统100提供语音通知，识别触发器和/或指示车辆200正制动或正要开始制动。

[0393] 在处理单元110识别触发器之后，在步骤3430，处理单元110可根据制动轮廓来致使车辆200减速和/或停止。在一些实施方式中，制动轮廓可包括一定数目的轮廓区段。制动轮廓的不同轮廓区段可由时间周期和/或车辆速度限定。例如，处理单元110可依据持续预定时间段或直到车辆200的速度到达目标值的第一轮廓区段来执行车辆200的控制。目标值可以是预定速度，或者可以是车辆200过去在执行制动轮廓之前的行进的速度的百分比。为了实施制动轮廓，处理单元110可向车辆200的节流系统220、制动系统230和/或转向系统240中的一个或多个发送电子信号。进一步，一个或多个致动器可控制节流系统220、制动系统230和/或转向系统240。

[0394] 如上面讨论的，根据一些实施方式，制动轮廓的轮廓区段可对应于减速率。例如，制动轮廓可包括多个区段，并且第一区段可与第一减速率关联。制动轮廓可进一步包括与第二减速率关联的第二区段。第二减速率可小于第一减速率。

[0395] 进一步，如上面讨论的，根据一些实施方式，制动轮廓可包括第三轮廓区段，其中制动水平的降低是车辆200与产生触发器的潜在因素(例如，对象3220、交通灯3230、交通标志3240)之间的距离的函数。例如，制动轮廓可包括在车辆200接近目标停止地点3250时降低制动水平。进一步，到达目标停止地点3250可伴随着车辆200的制动压力的释放。

[0396] 出于图示的目的，前面的描述已呈现。并非穷尽且并不限于公开的具体形式或实施方式。考虑到说明书和公开的实施方式的实践，修改和改写对本领域技术人员将是明显的。另外，虽然公开的实施方式的各方面描述为存储在存储器中，但是本领域技术人员将了解到，这些方面也可以存储在其它类型的诸如二级存储装置的计算机可读介质上，例如，硬盘或CD ROM，或其它形式的RAM或ROM、USB介质、DVD、蓝光或其它光驱介质。

[0397] 基于书面描述和公开方法的计算机程序在有经验的开发人员的技术内。各种程序或程序模块可以使用本领域技术人员公知的任何技术来创建或者可以结合现有软件设计。例如，程序段或程序模块可以设计或借助.Net Framework、.Net Compact Framework(及相关语言，诸如Visual Basic、C，等等)、Java、C++、Objective-C、HTML、HTML/AJAX组合、XML或包括Java小程序的HTML。

[0398] 而且，虽然说明性实施方式已在本文中描述，具有等同元件、修改、省略、组合(例如，跨各实施方式的方面)、改写和/或变更的任何及所有实施方式的范围将基于本公开被本领域技术人员了解。权利要求书中的限制应解译宽泛地基于权利要求书中采用的语言，并不限于本说明书或本申请审查期间描述的示例。示例应被解释为非排他性的。此外，公开方法中的步骤能以任何方式修改，包括重新排序各步骤和/或插入或删除步骤。因此，预期说明书和示例考虑为仅是说明性的，真正的范围和精神由以下权利要求书及其完整的等同范围指示。

[0399] 相关申请的交叉引用

IPR2025-00943
Tesla EX1002 Page 1310

[0400] 本申请要求以下的优先权：2013年12月4日提交的美国临时专利申请号61/911,490；2014年5月14日提交的美国临时专利申请号61/993,084；2014年5月14日提交的美国临时专利申请号61/993,111；2014年6月23日提交的美国临时专利申请62/015,524；2014年7月9日提交的美国临时专利申请62/022,221；2014年8月21日提交的美国临时专利申请号62/040,224；以及2014年8月21日提交的美国临时专利申请62/040,269。所有前述申请的整体内容都通过引用并入本文。

IPR2025-00943
Tesla EX1002 Page 1311

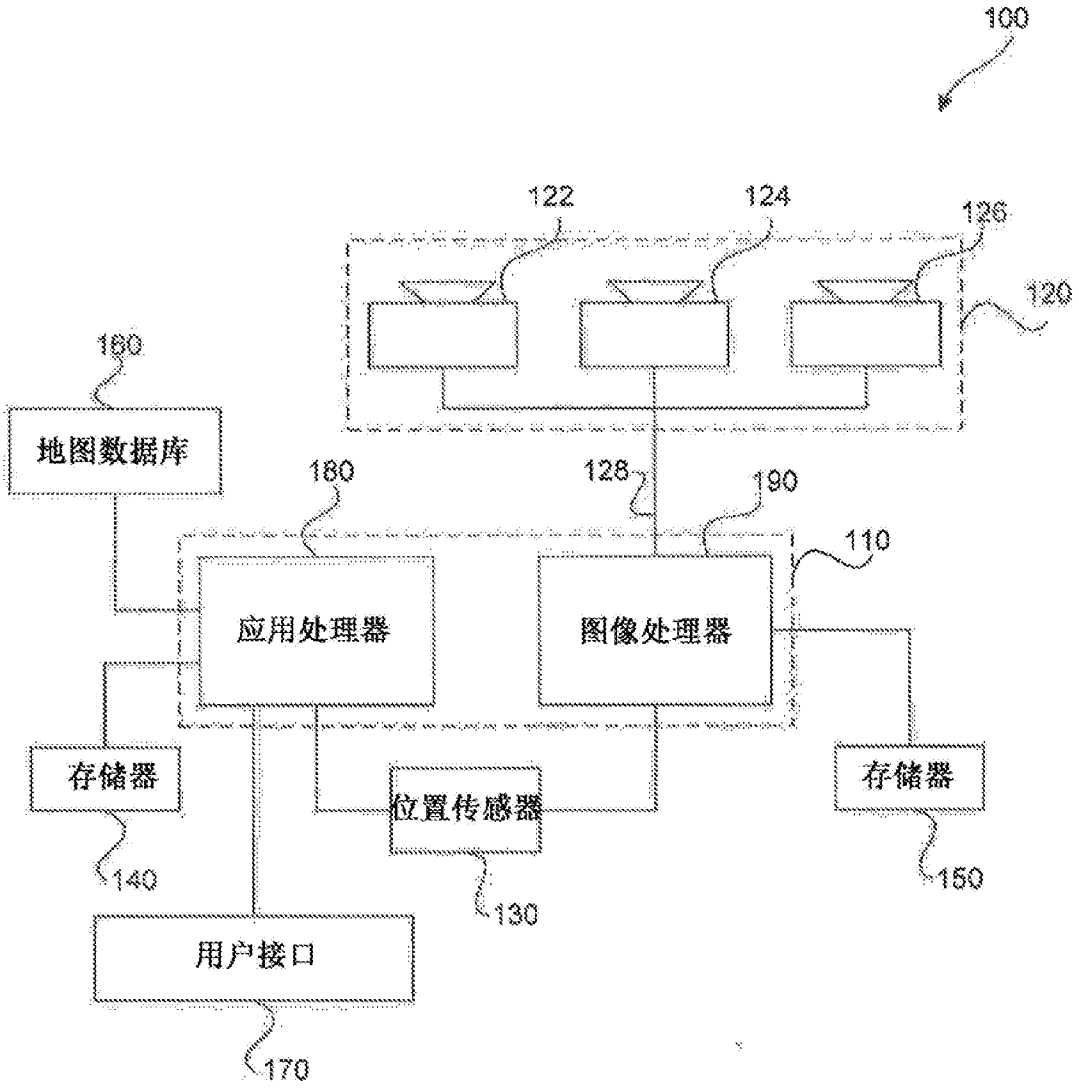


图1

IPR2025-00943
Tesla EX1002 Page 1312

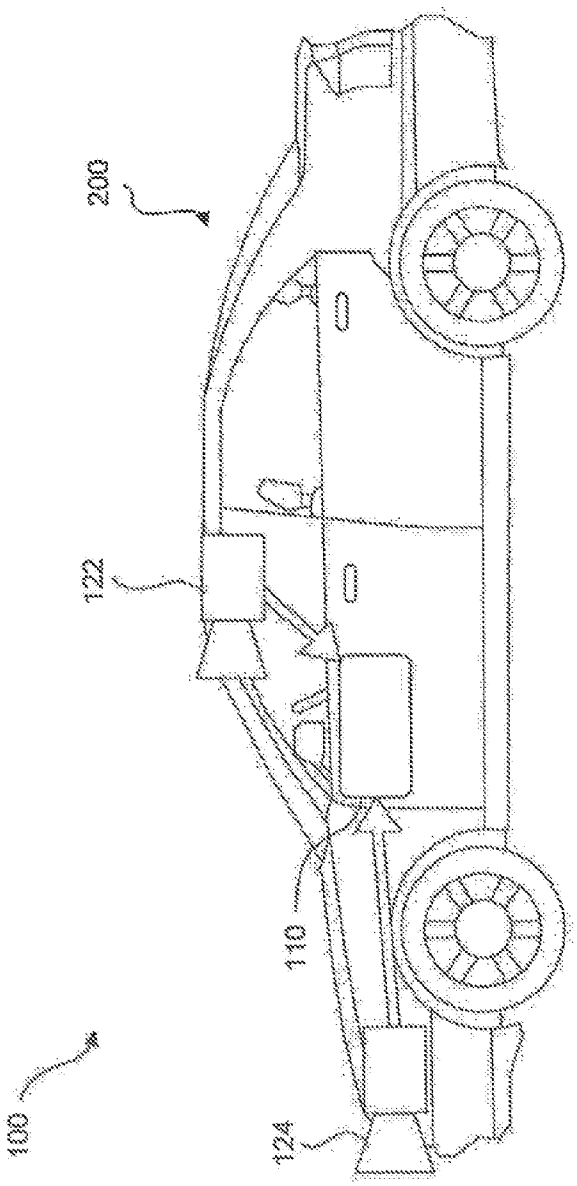


图2A

IPR2025-00943
Tesla EX1002 Page 1313

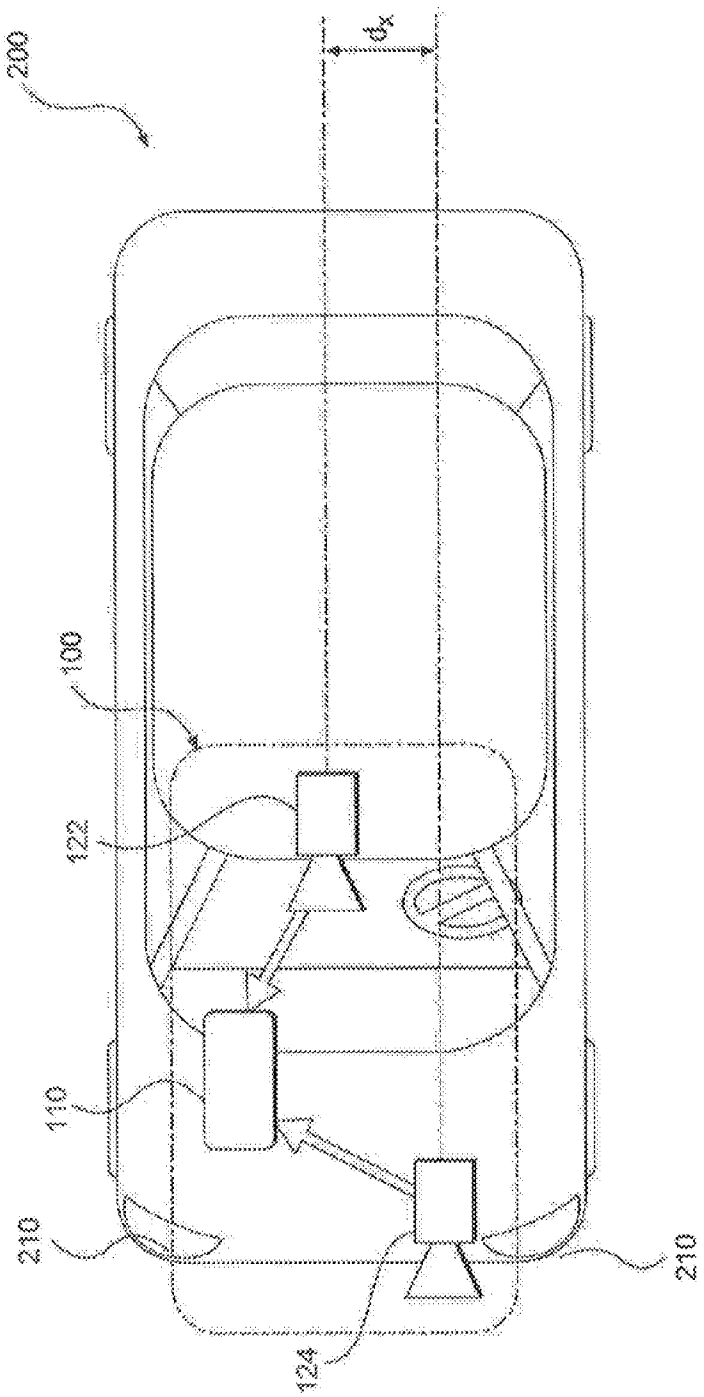


图2B

IPR2025-00943
Tesla EX1002 Page 1314

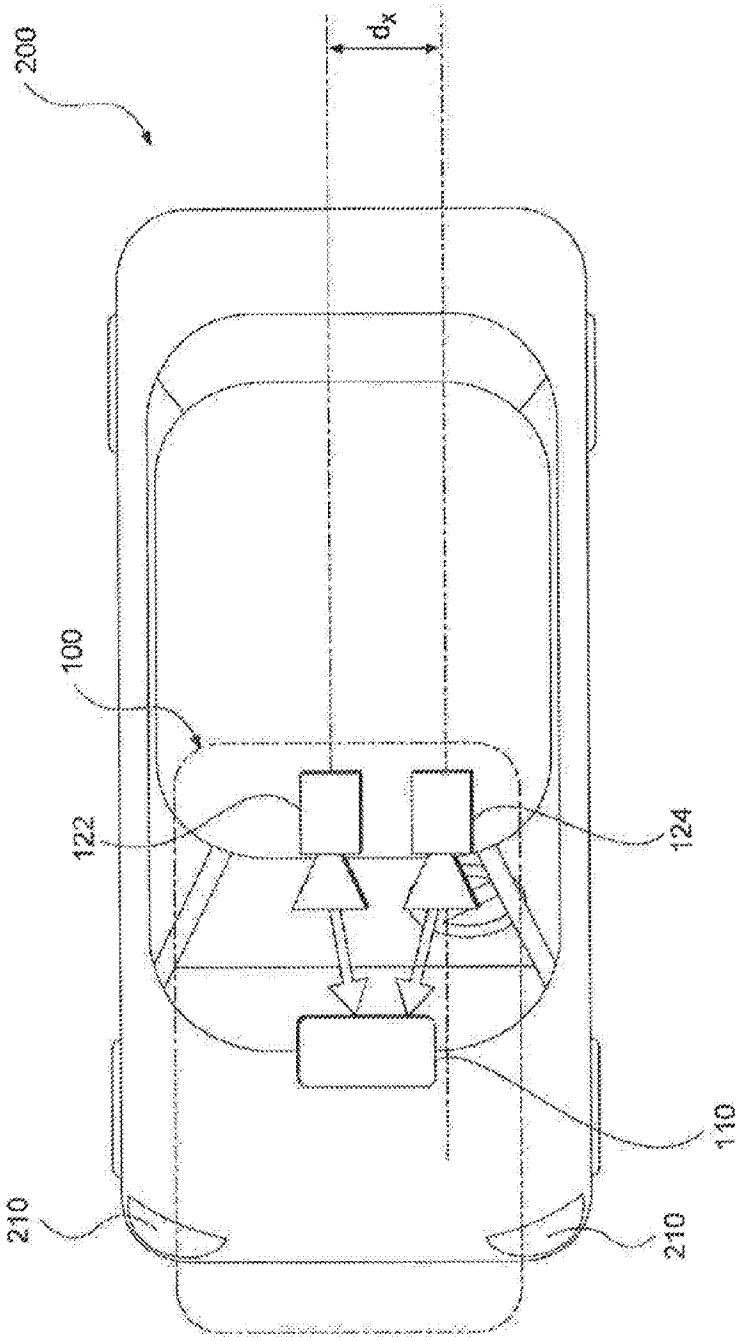


图2C

IPR2025-00943
Tesla EX1002 Page 1315

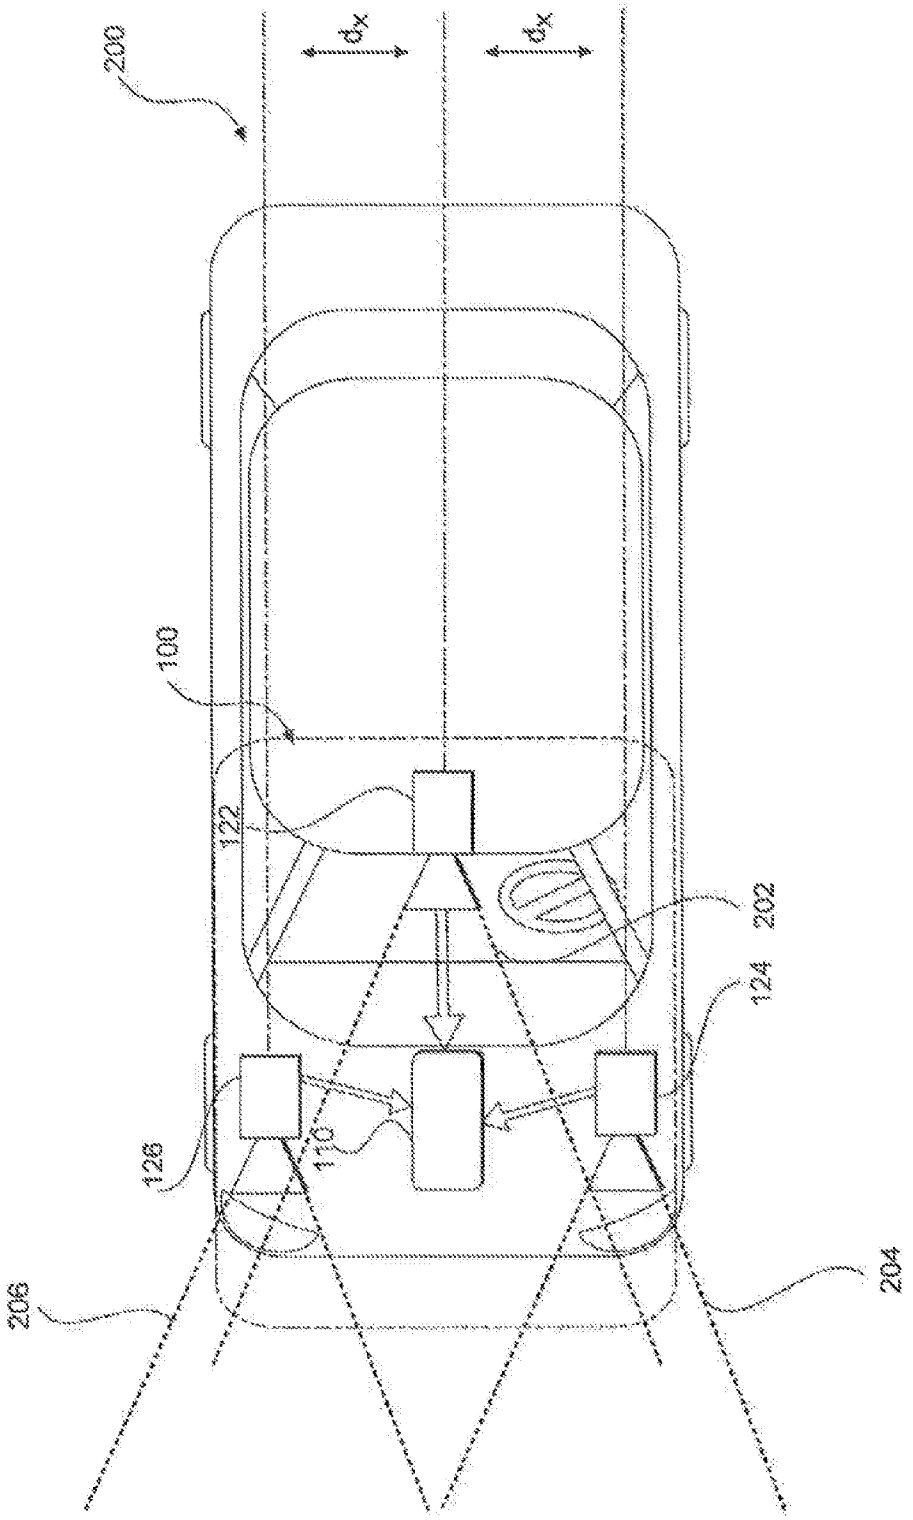


图2D

IPR2025-00943
Tesla EX1002 Page 1316

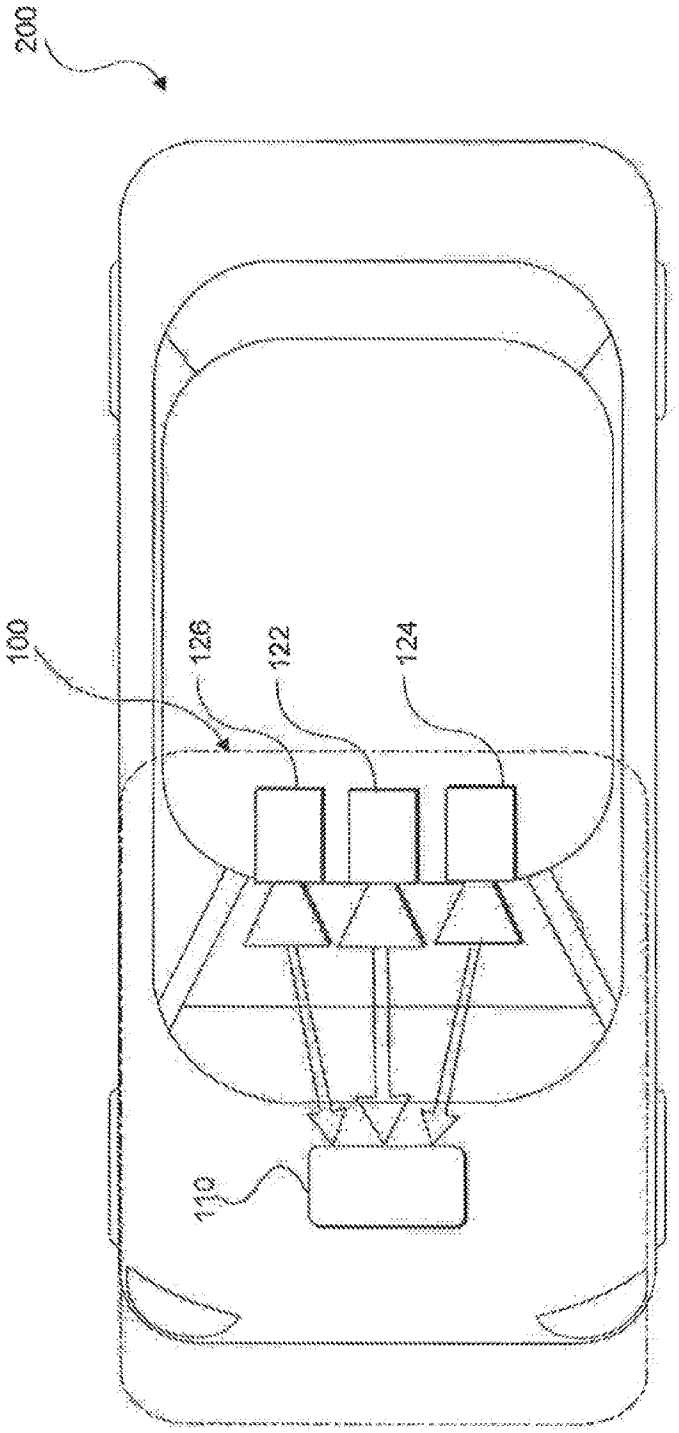


图2E

IPR2025-00943
Tesla EX1002 Page 1317

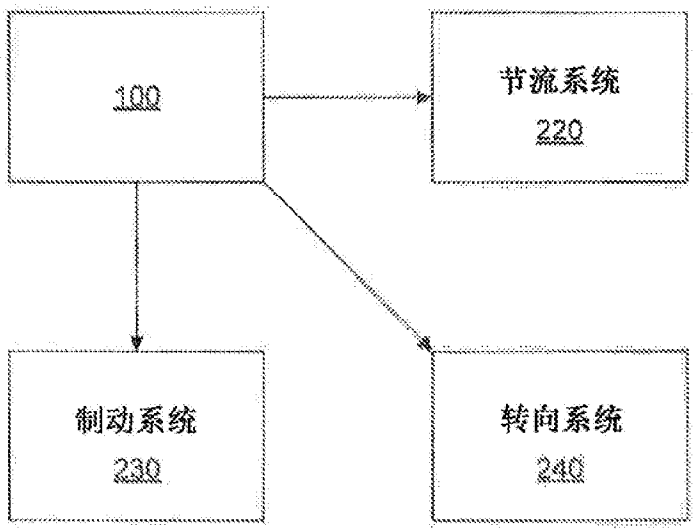


图2F

IPR2025-00943
Tesla EX1002 Page 1318

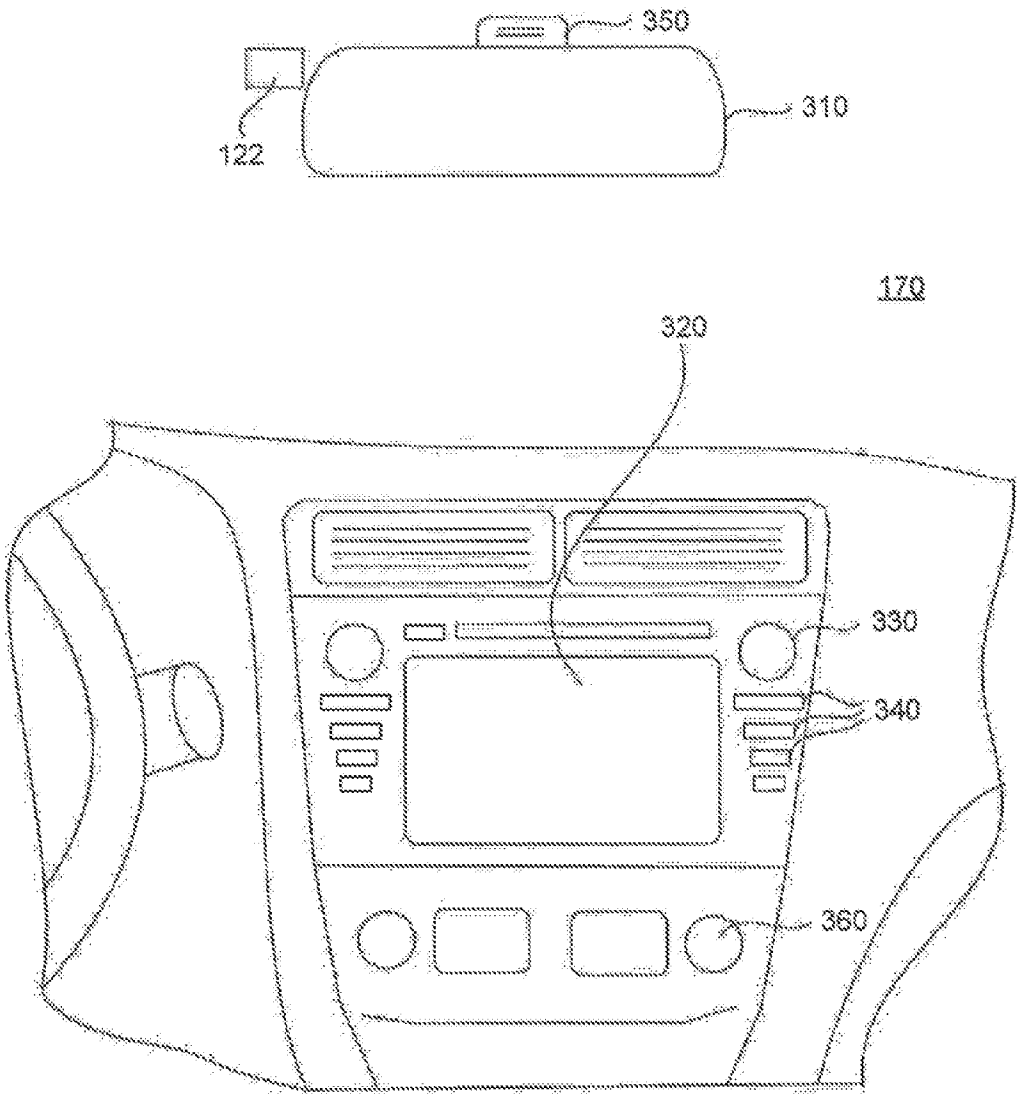


图3A

IPR2025-00943
Tesla EX1002 Page 1319

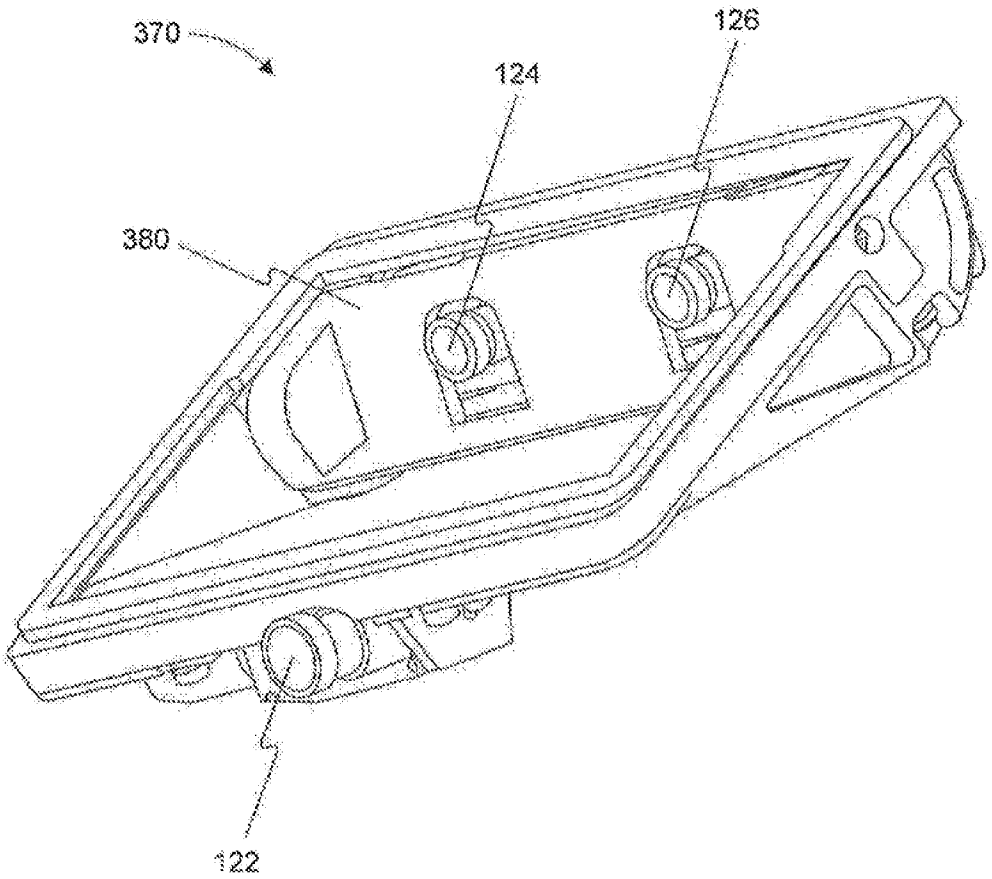


图3B

IPR2025-00943
Tesla EX1002 Page 1320

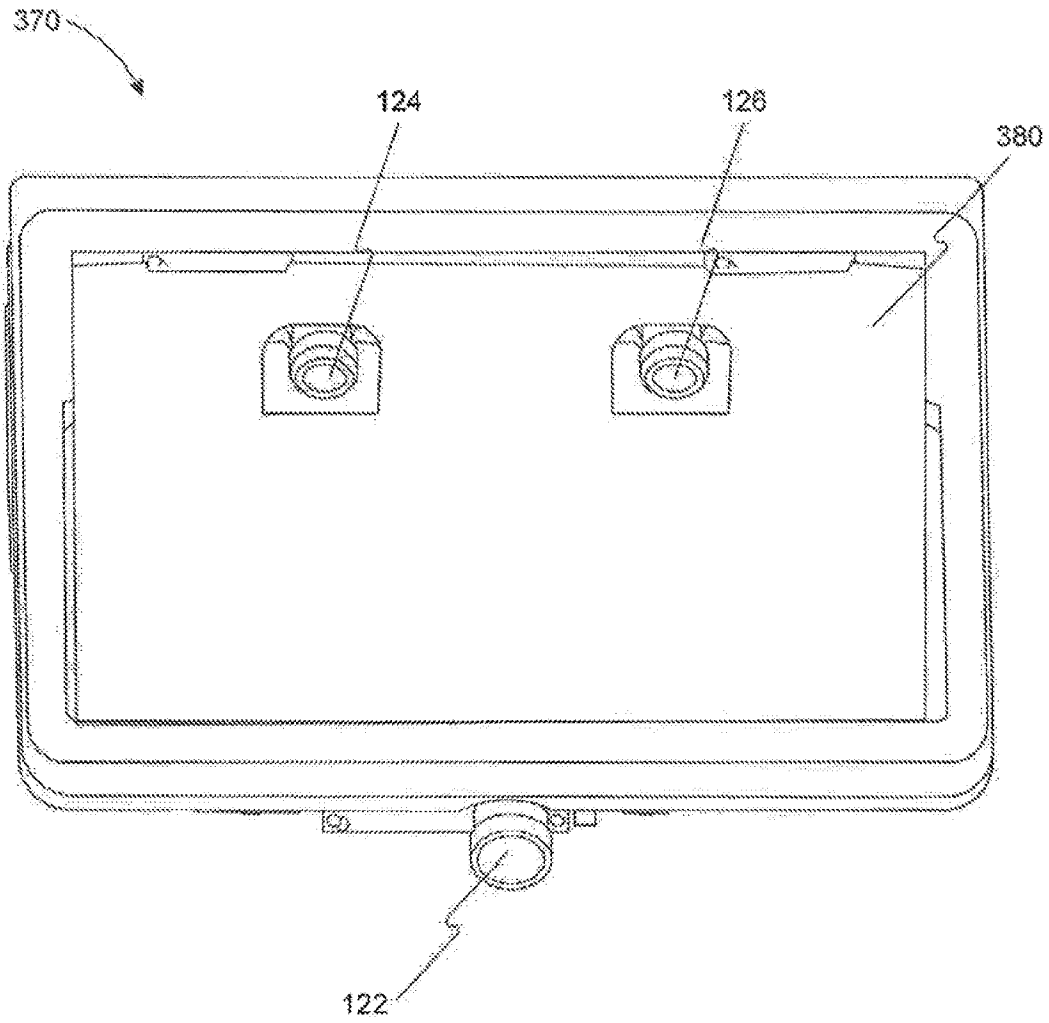


图3C

IPR2025-00943
Tesla EX1002 Page 1321

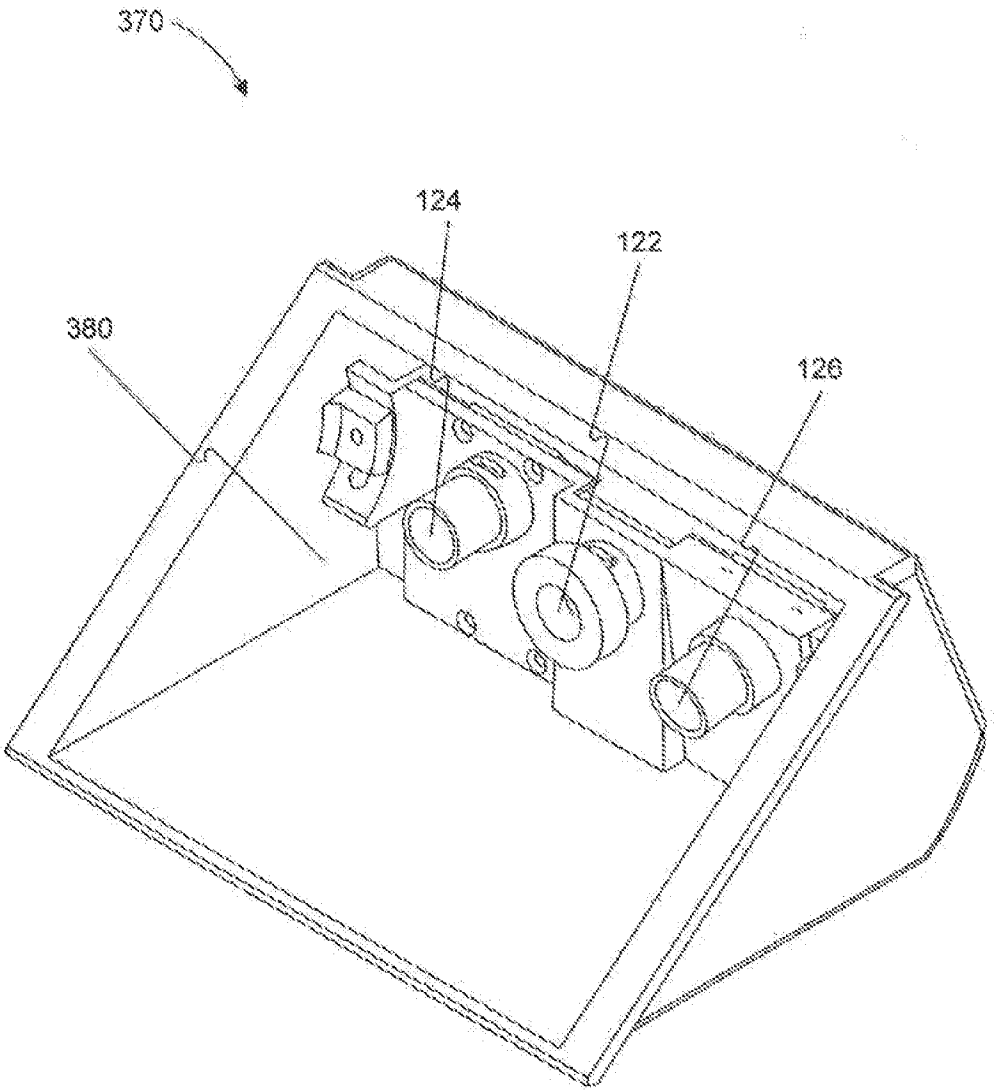


图3D

IPR2025-00943
Tesla EX1002 Page 1322

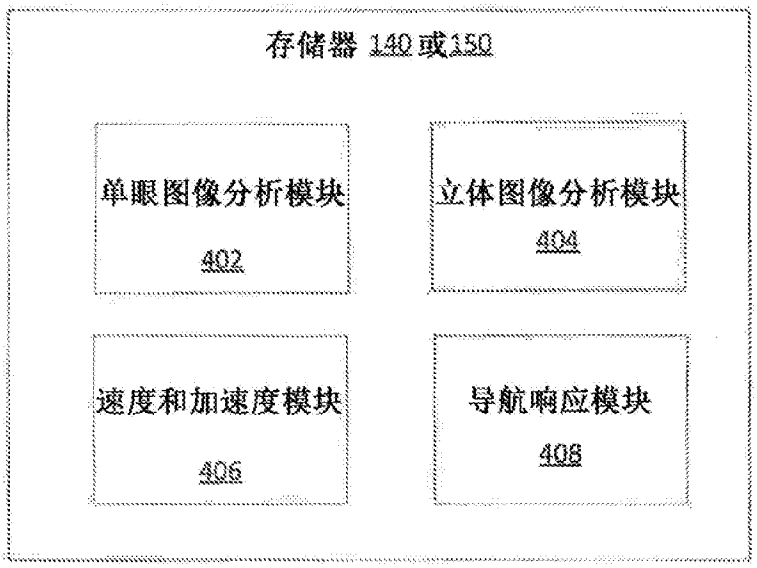


图4

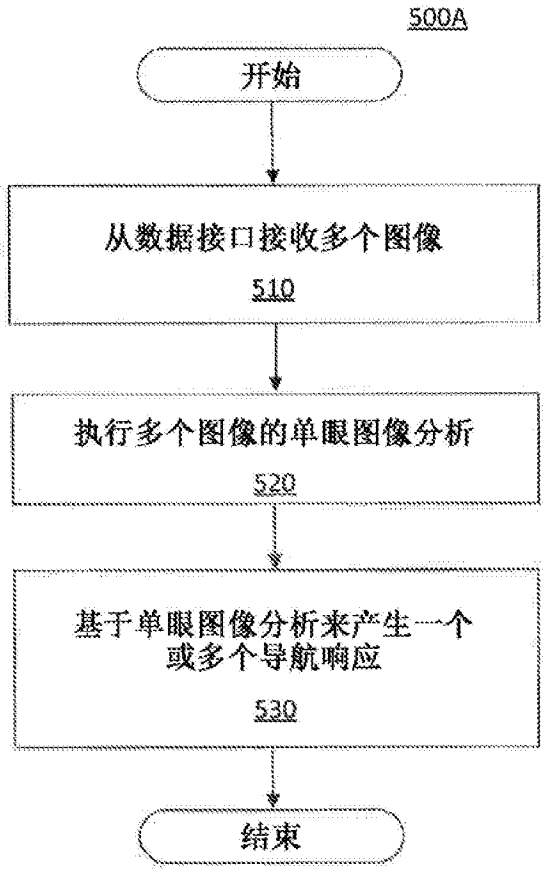


图5A

IPR2025-00943
Tesla EX1002 Page 1323

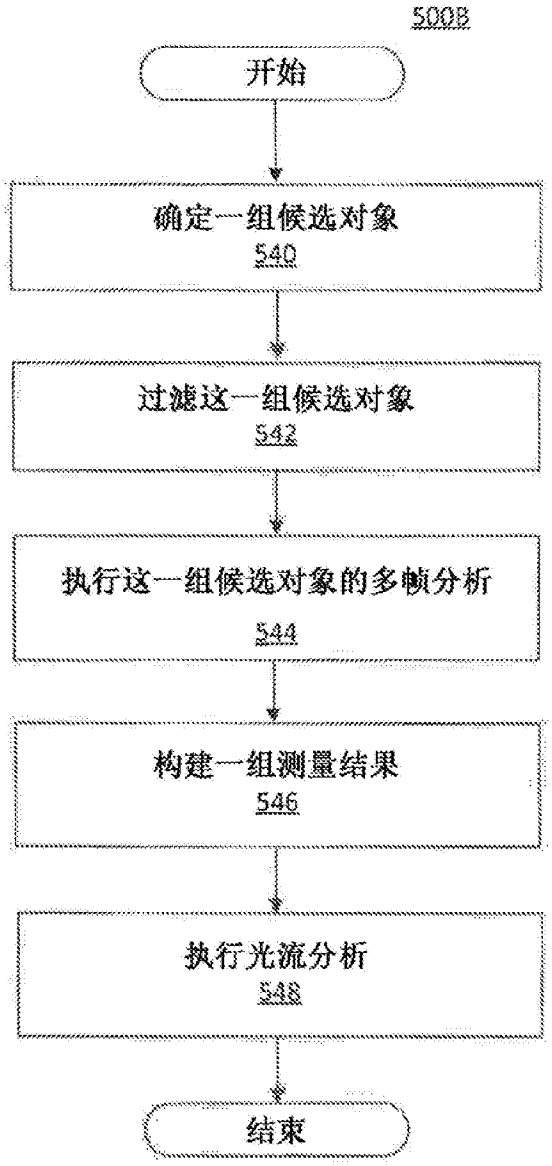


图5B

IPR2025-00943
Tesla EX1002 Page 1324

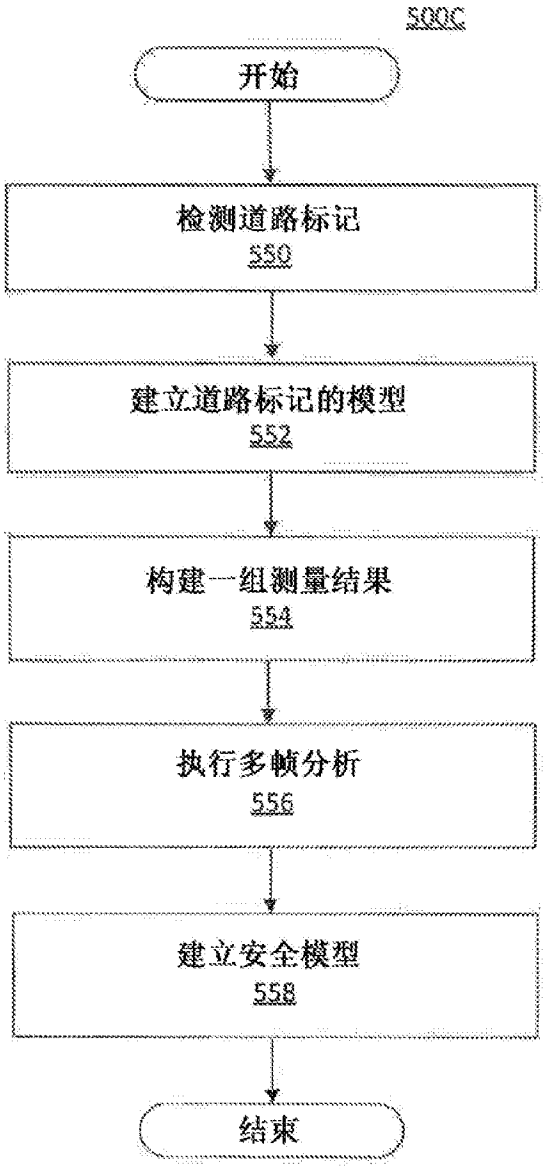


图5C

IPR2025-00943
Tesla EX1002 Page 1325

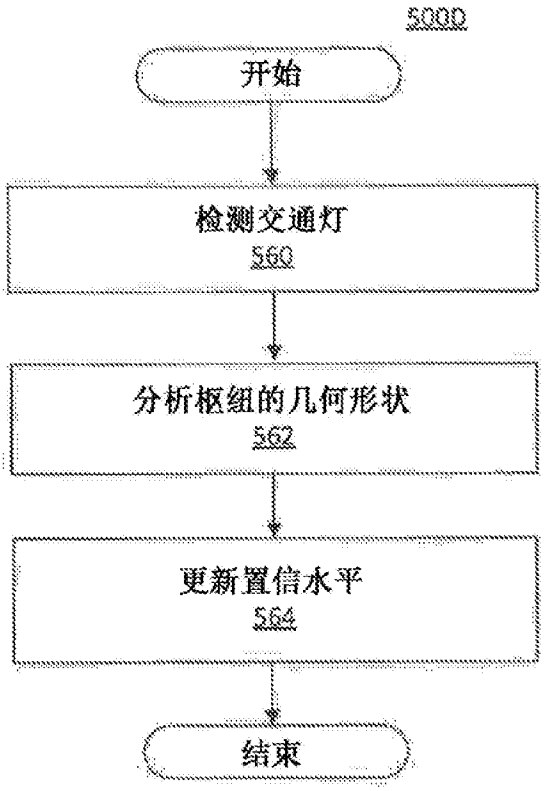


图5D

IPR2025-00943
Tesla EX1002 Page 1326

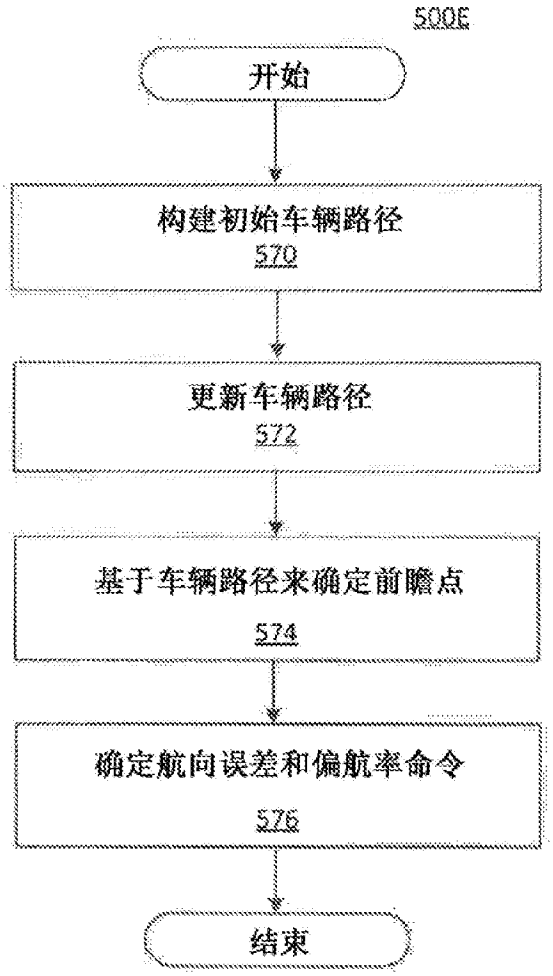


图5E

IPR2025-00943
Tesla EX1002 Page 1327

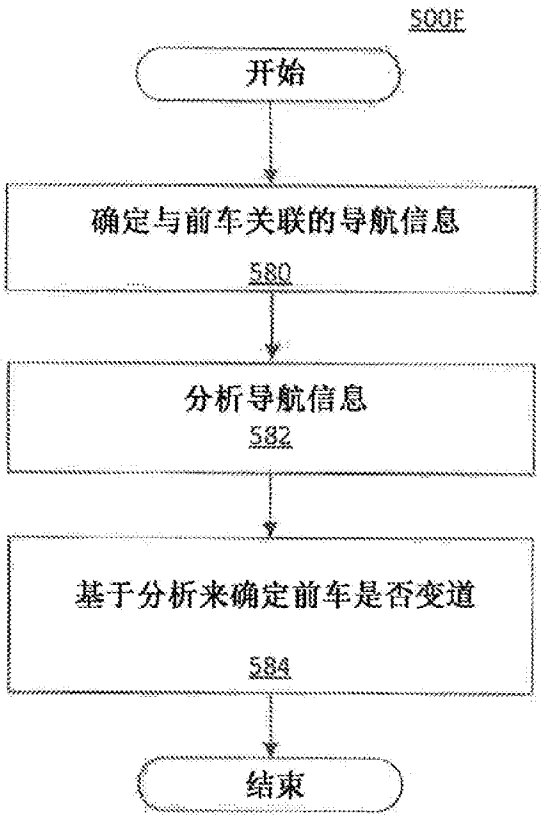


图5F

IPR2025-00943
Tesla EX1002 Page 1328

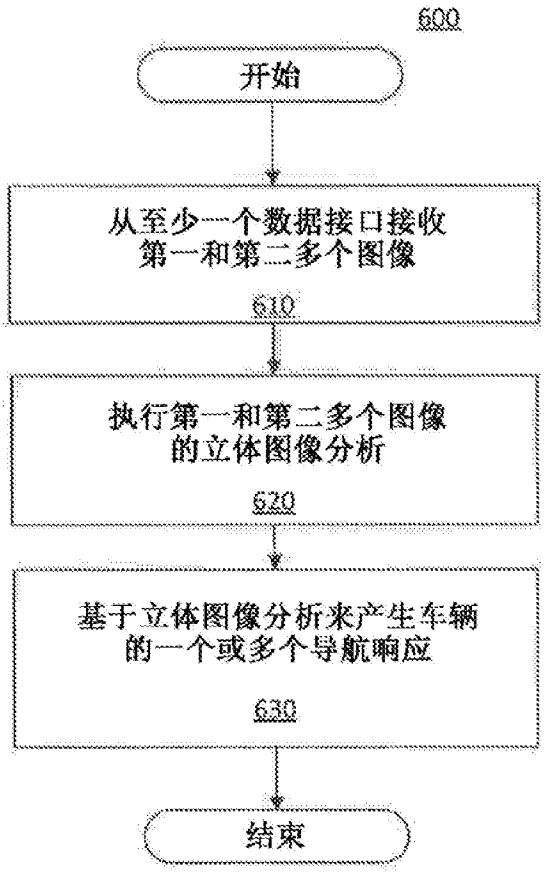


图6

IPR2025-00943
Tesla EX1002 Page 1329

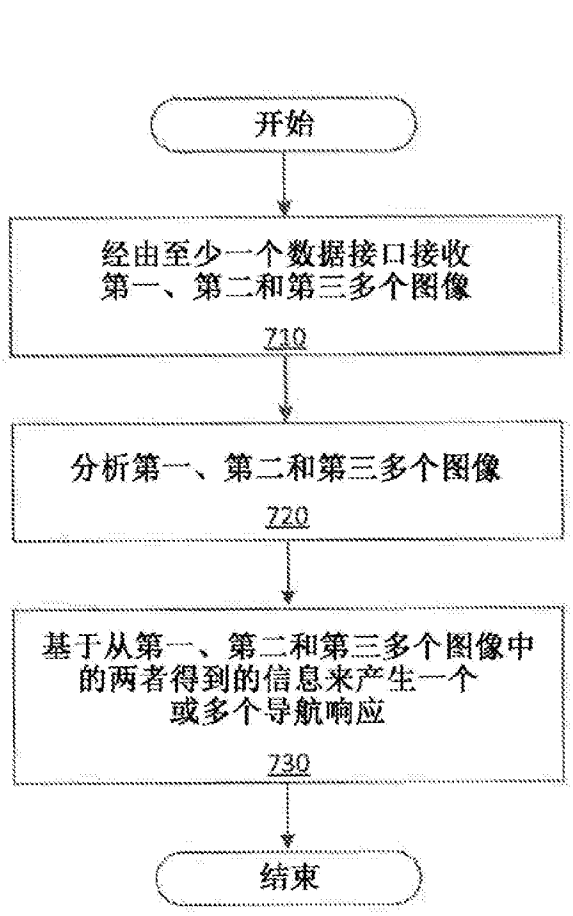


图7

IPR2025-00943
Tesla EX1002 Page 1330

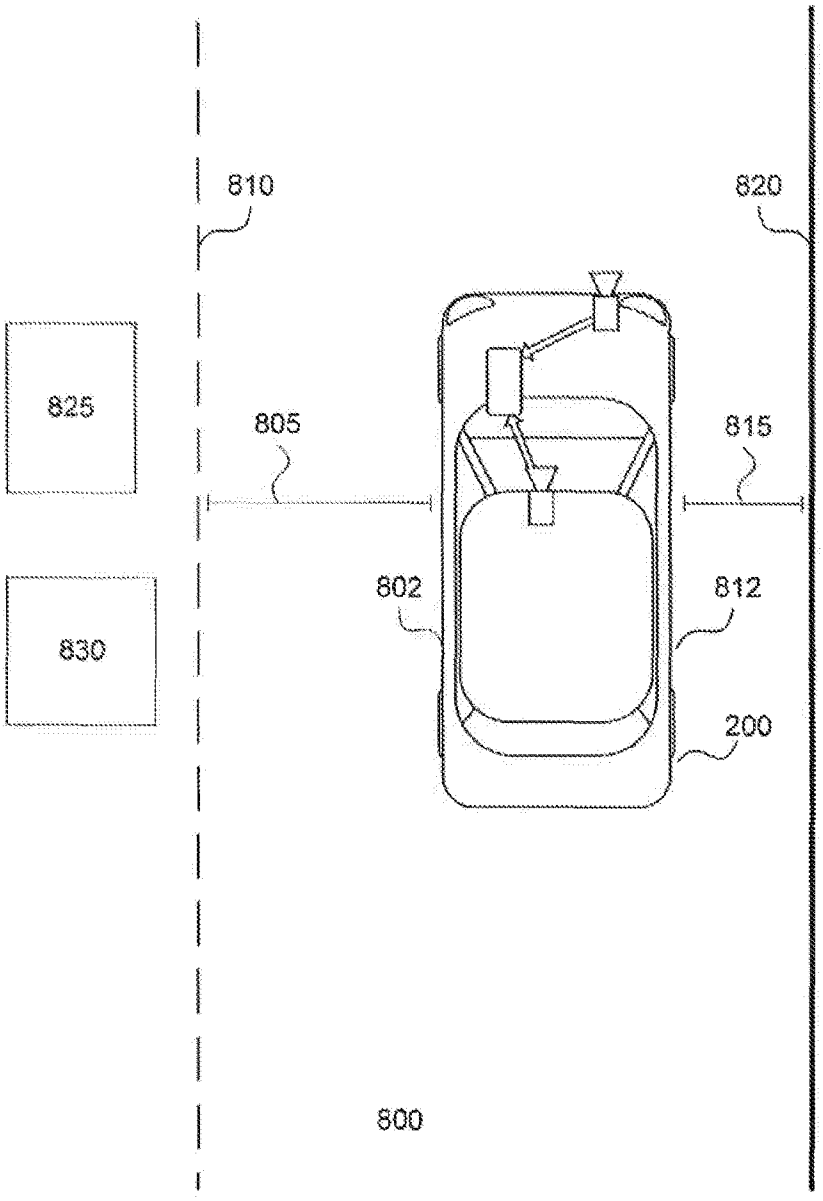


图8A

IPR2025-00943
Tesla EX1002 Page 1331

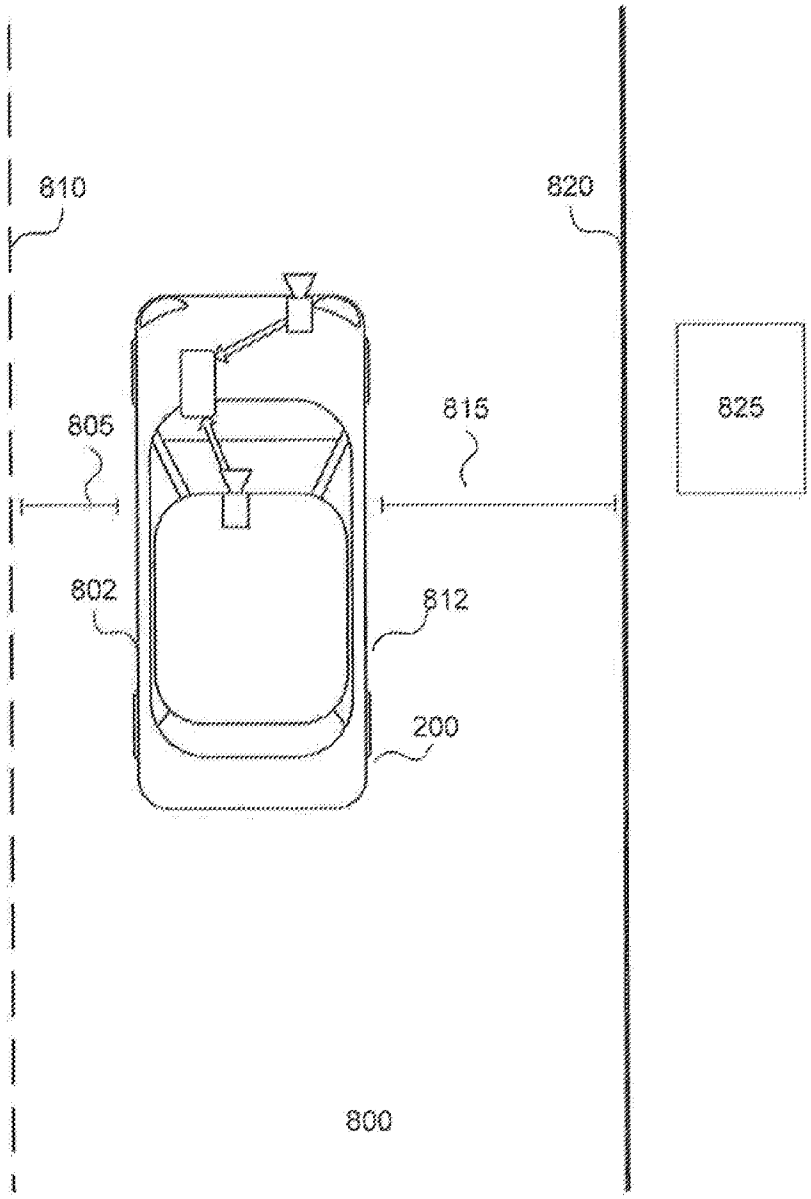


图8B

IPR2025-00943
Tesla EX1002 Page 1332

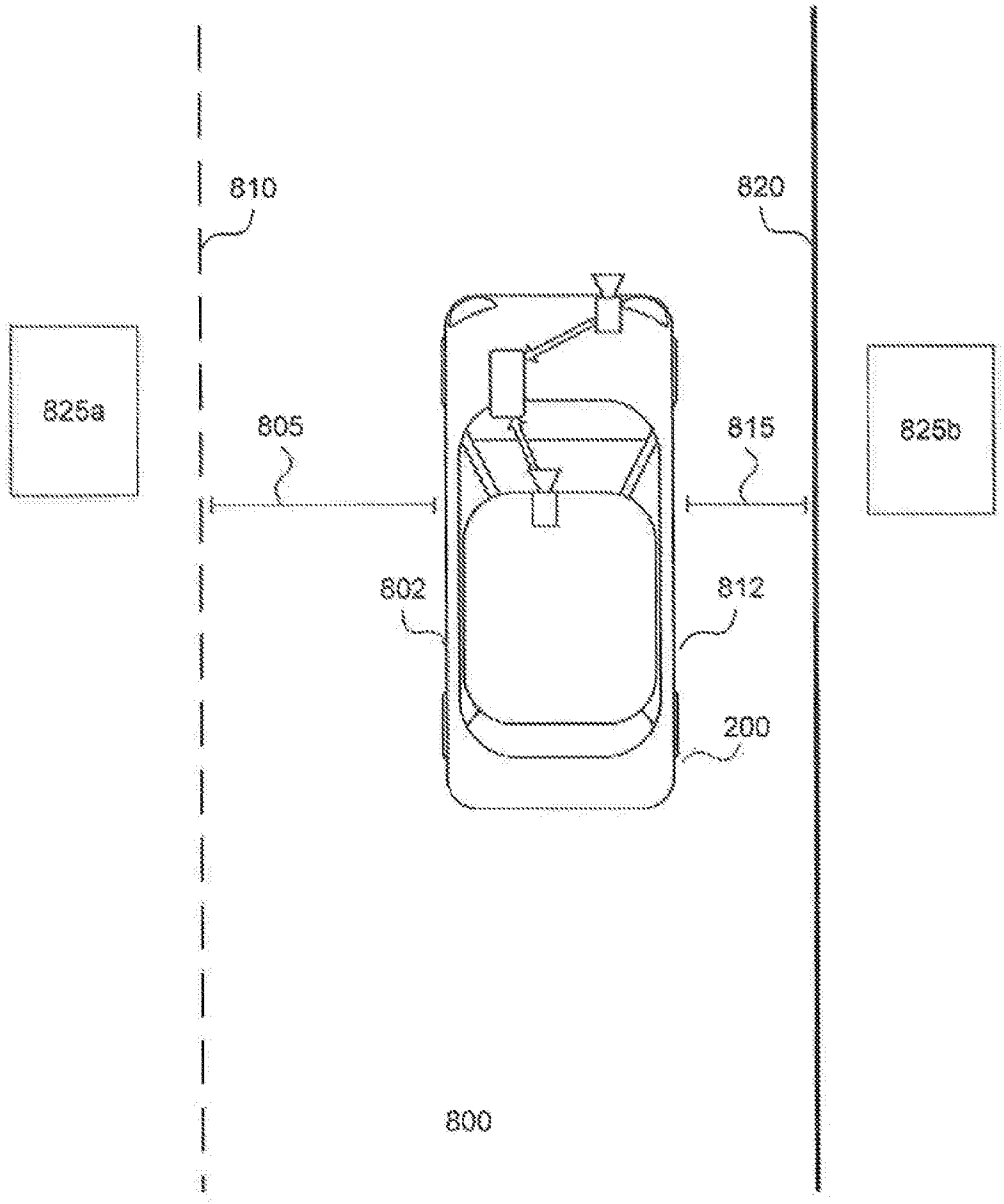


图8C

IPR2025-00943
Tesla EX1002 Page 1333

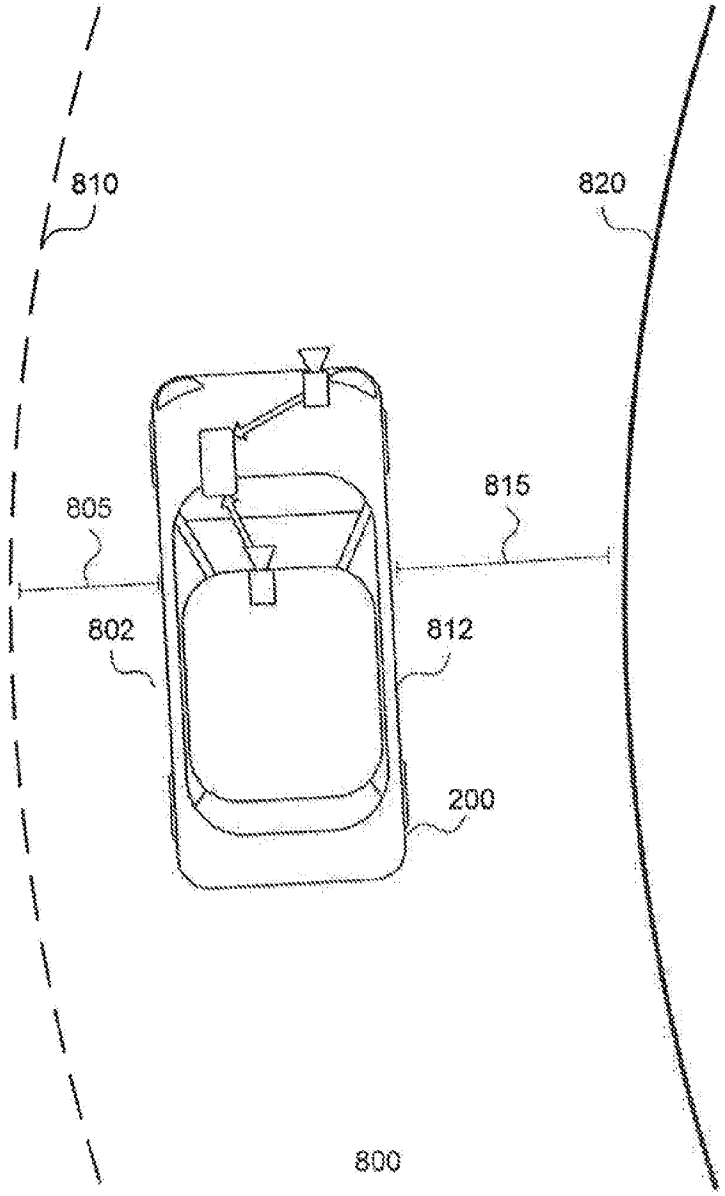


图8D

IPR2025-00943
Tesla EX1002 Page 1334

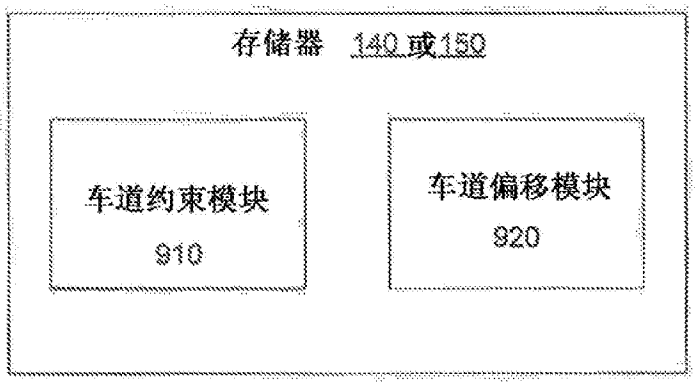


图9

IPR2025-00943
Tesla EX1002 Page 1335

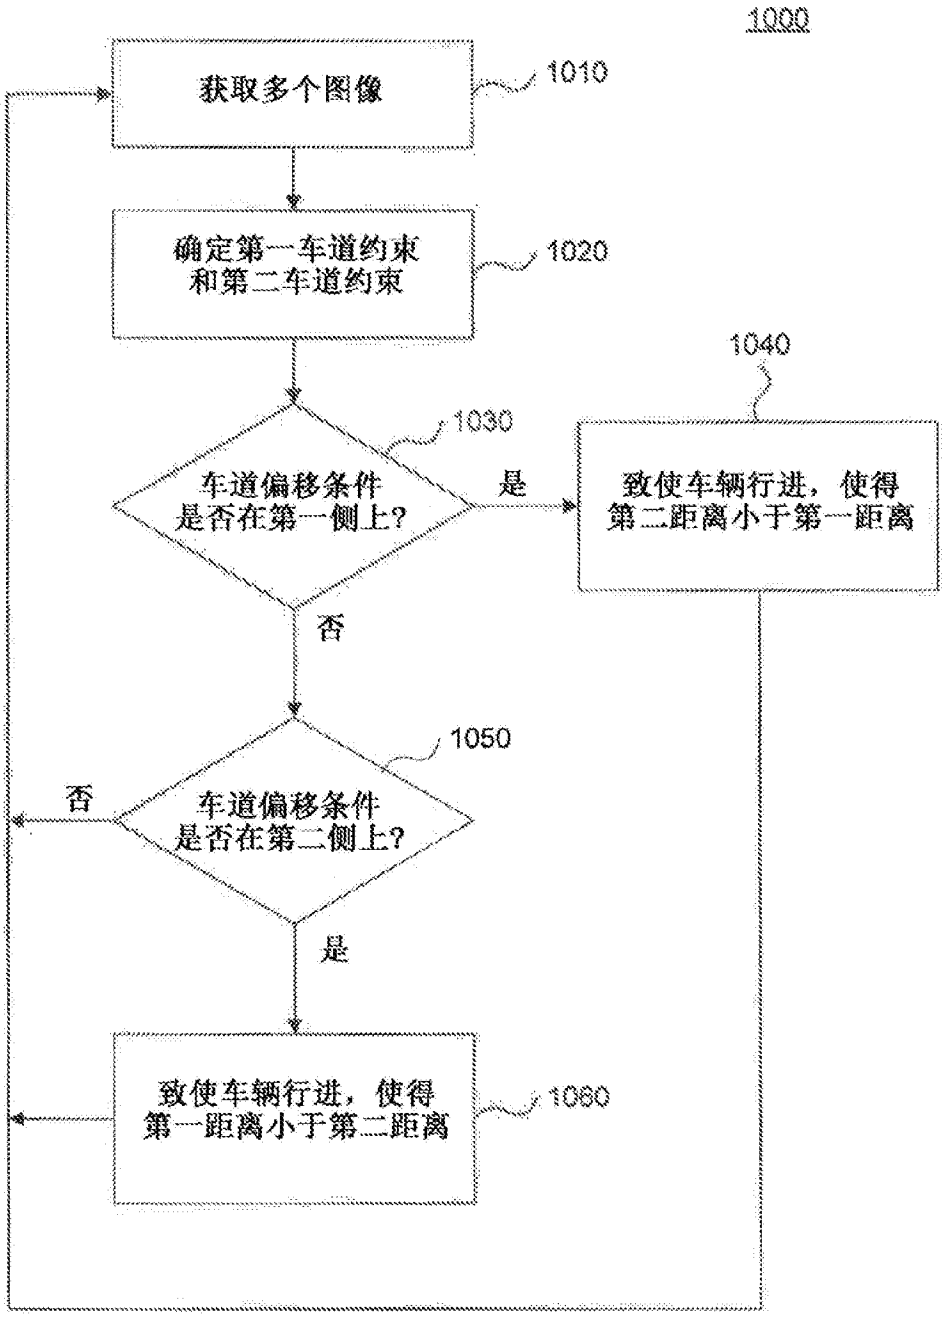


图10

IPR2025-00943
Tesla EX1002 Page 1336

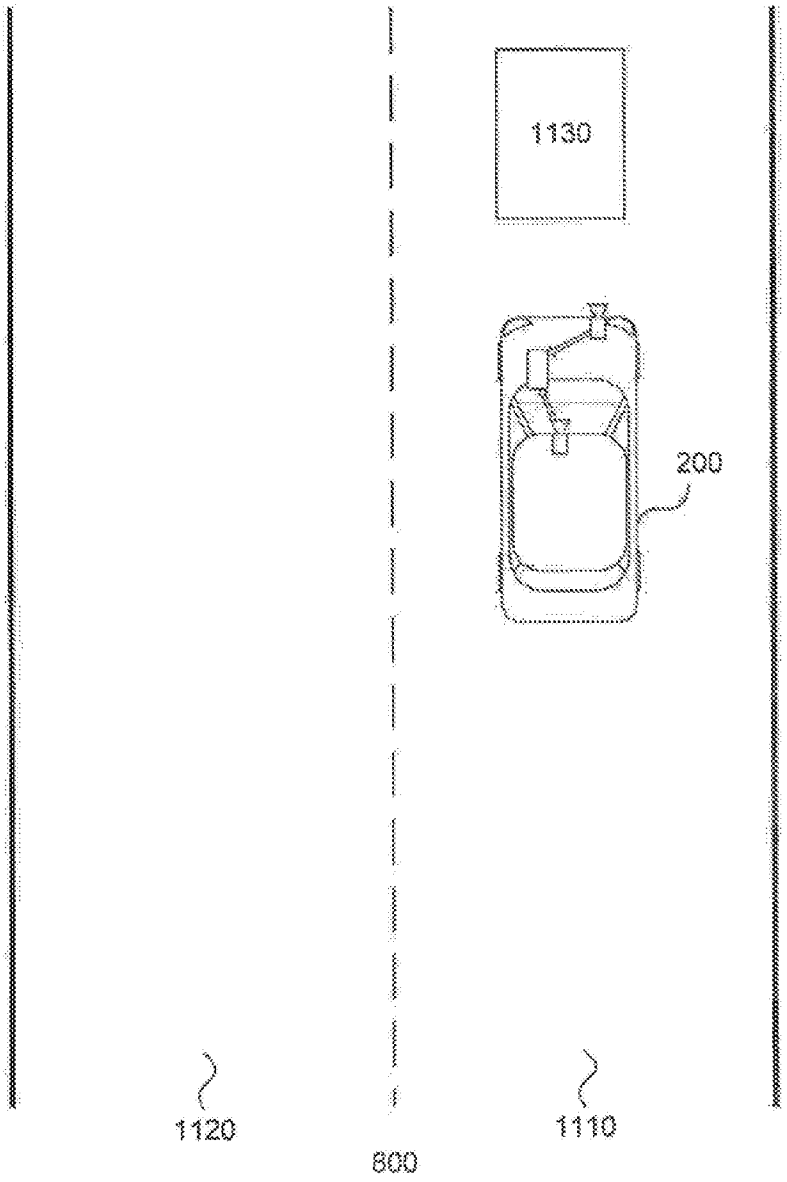


图11

IPR2025-00943
Tesla EX1002 Page 1337

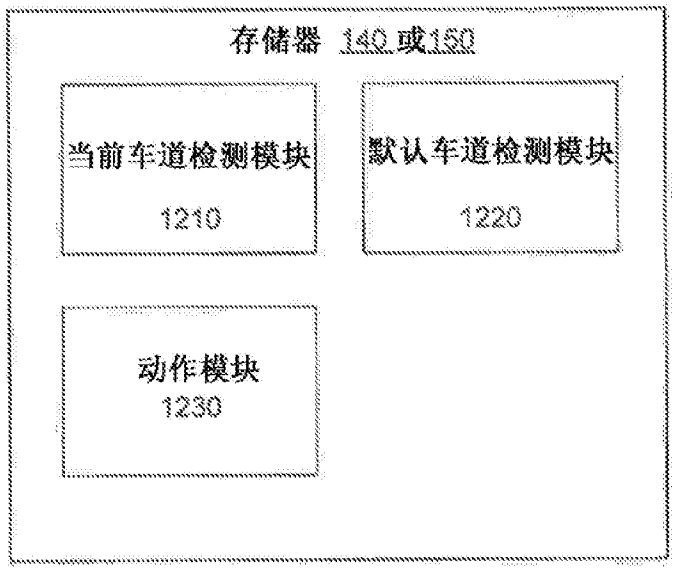


图12A

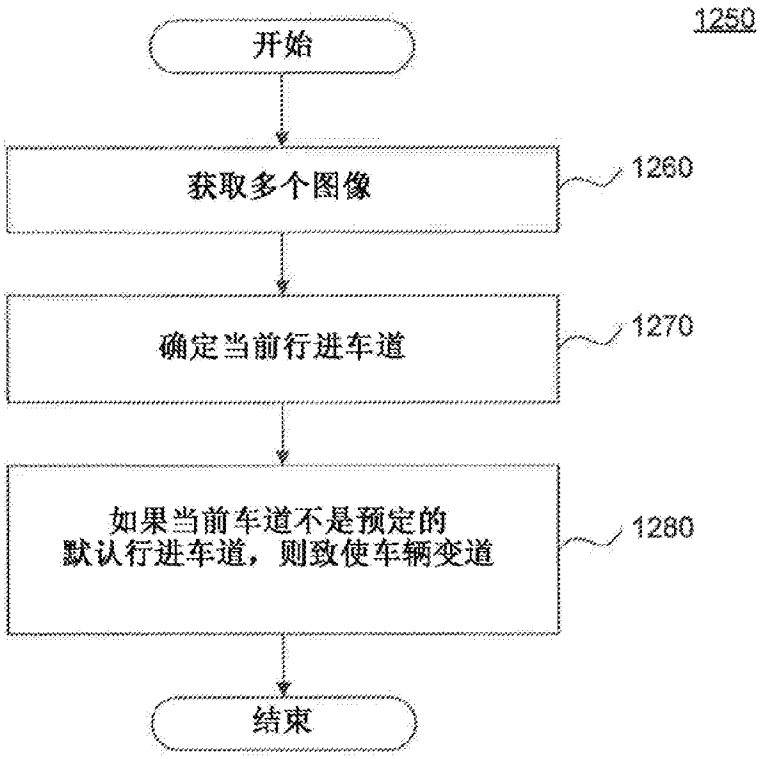


图12B

IPR2025-00943
Tesla EX1002 Page 1338

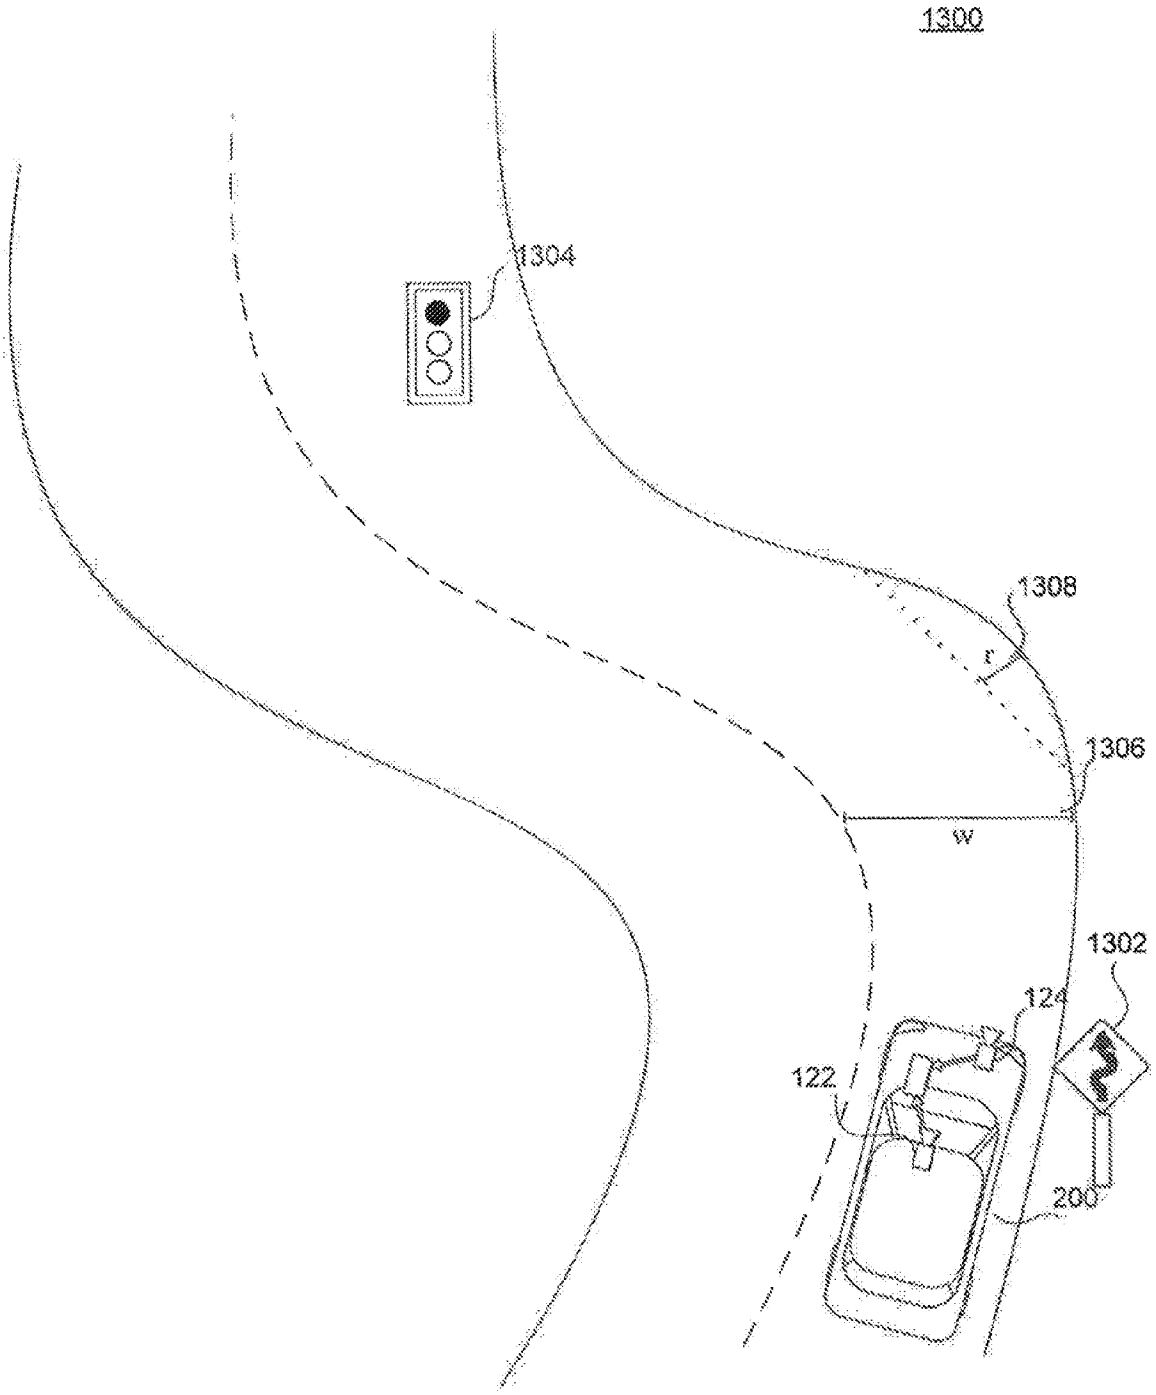


图13A

IPR2025-00943
Tesla EX1002 Page 1339

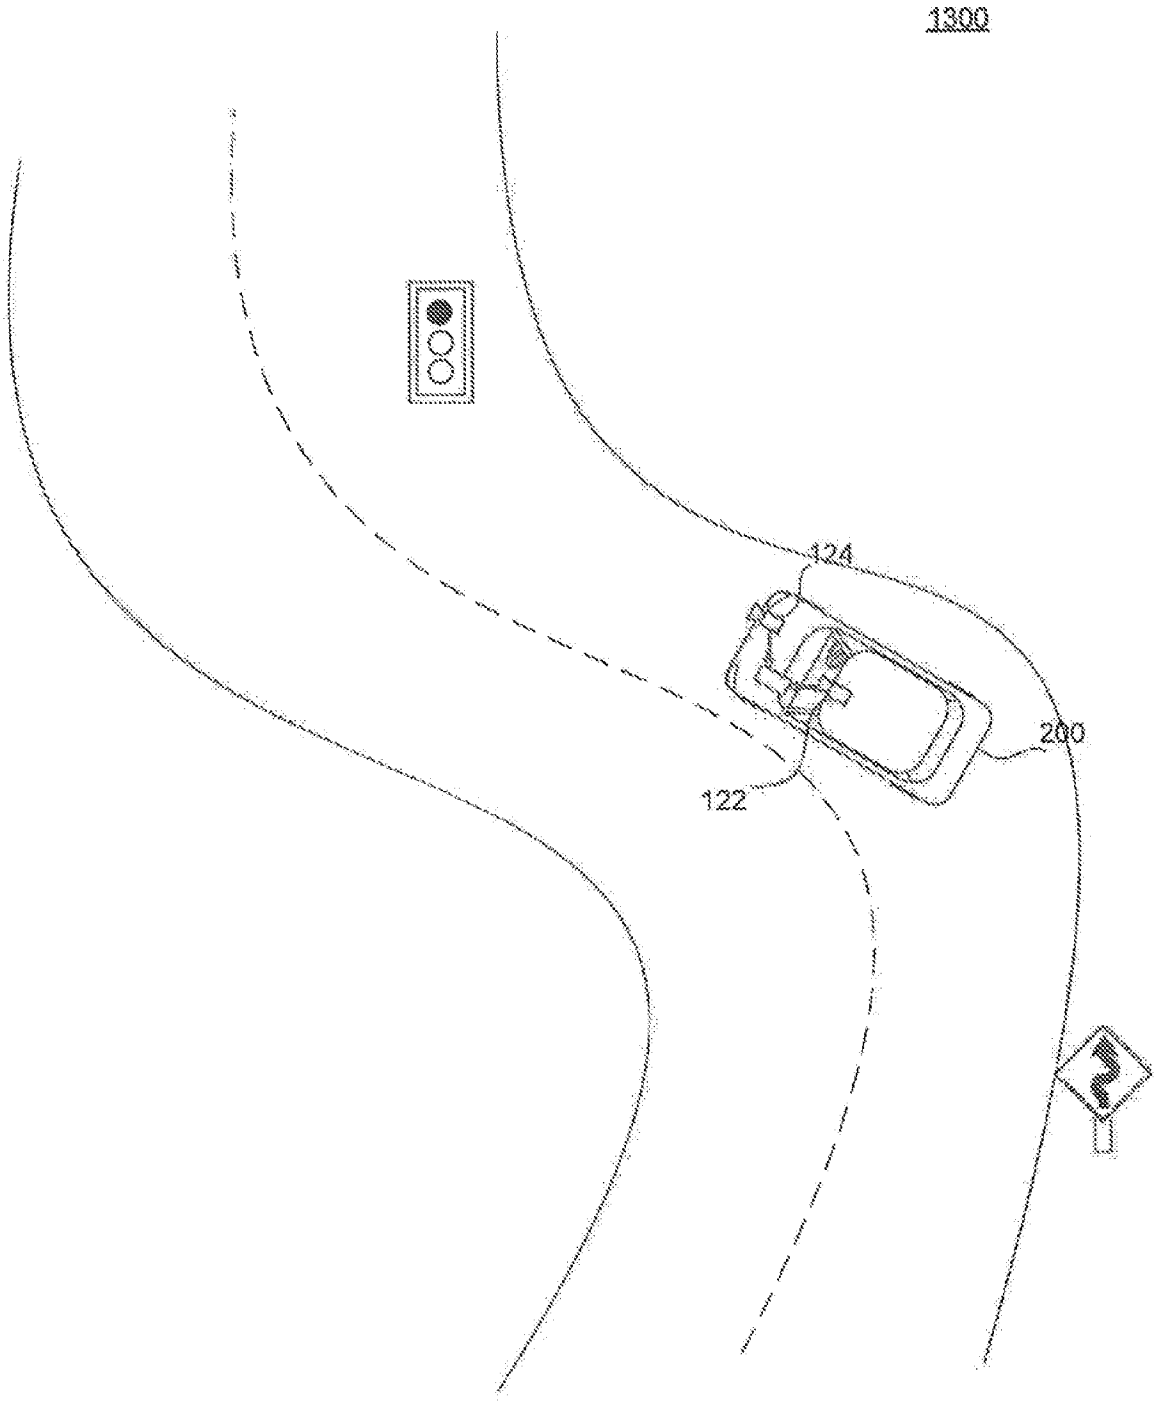


图13B

IPR2025-00943
Tesla EX1002 Page 1340

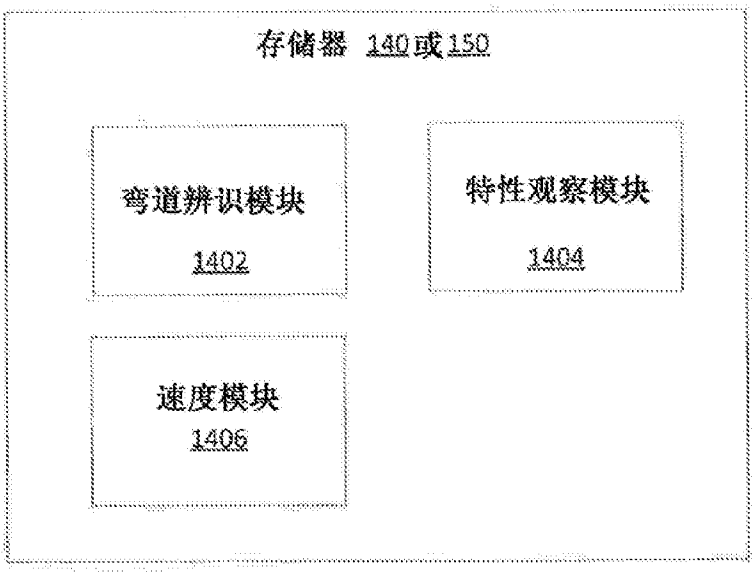


图14

IPR2025-00943
Tesla EX1002 Page 1341



图15

IPR2025-00943
Tesla EX1002 Page 1342

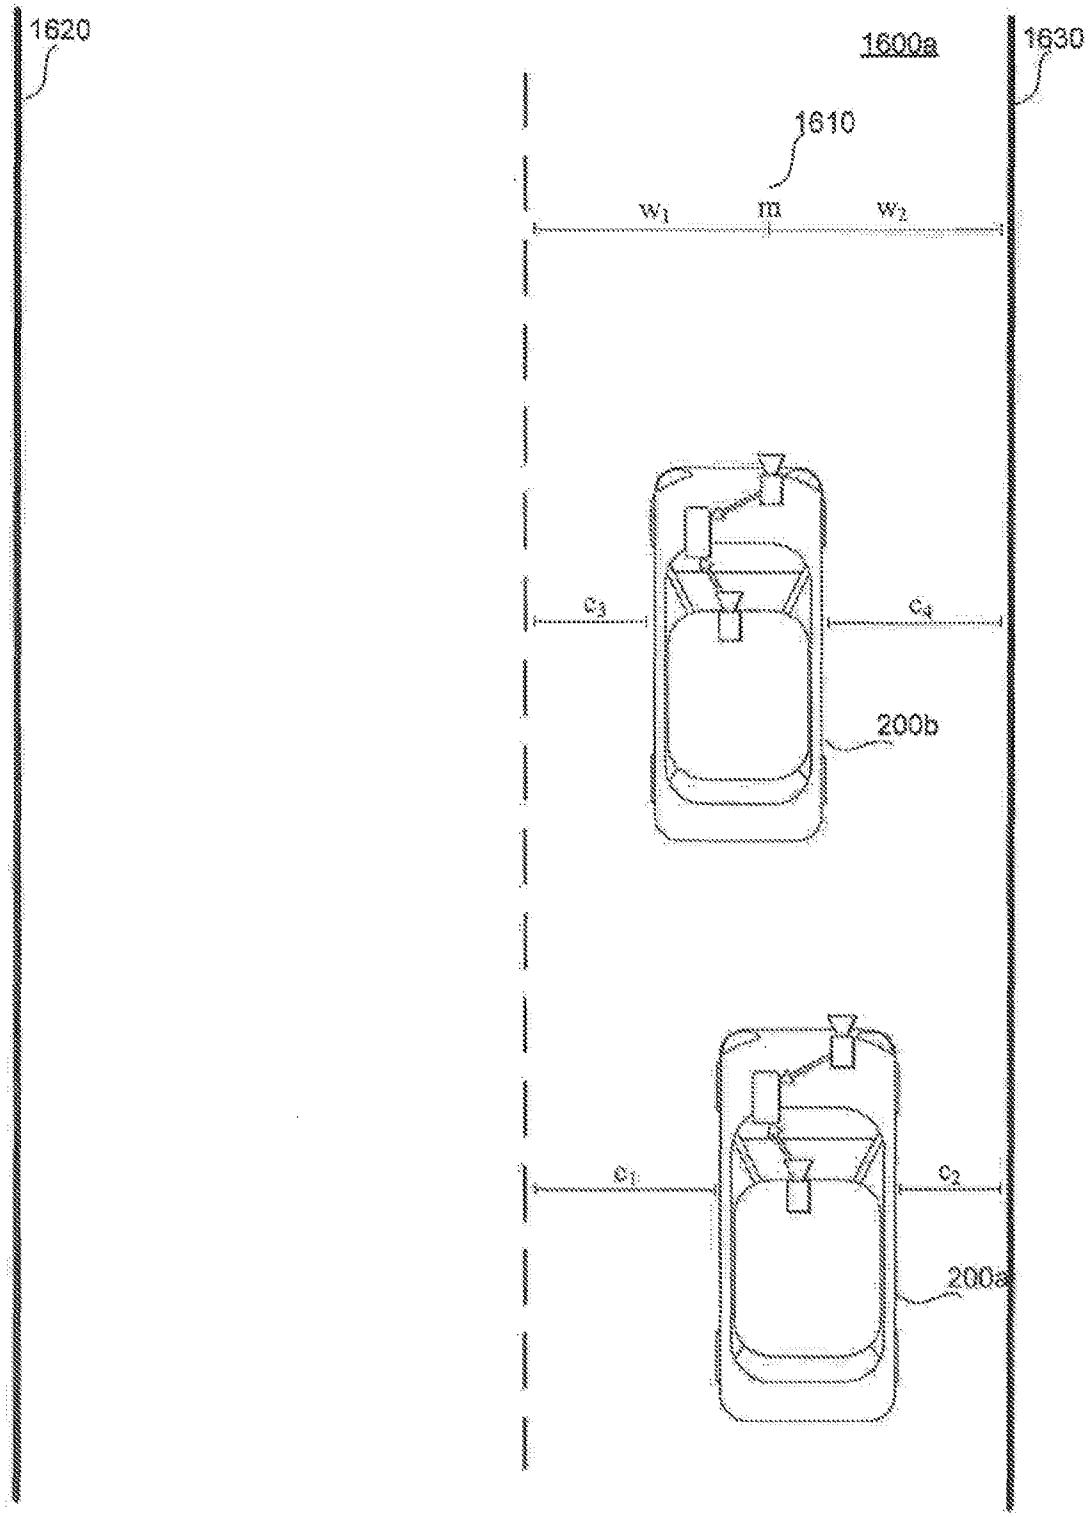


图16A

IPR2025-00943
Tesla EX1002 Page 1343

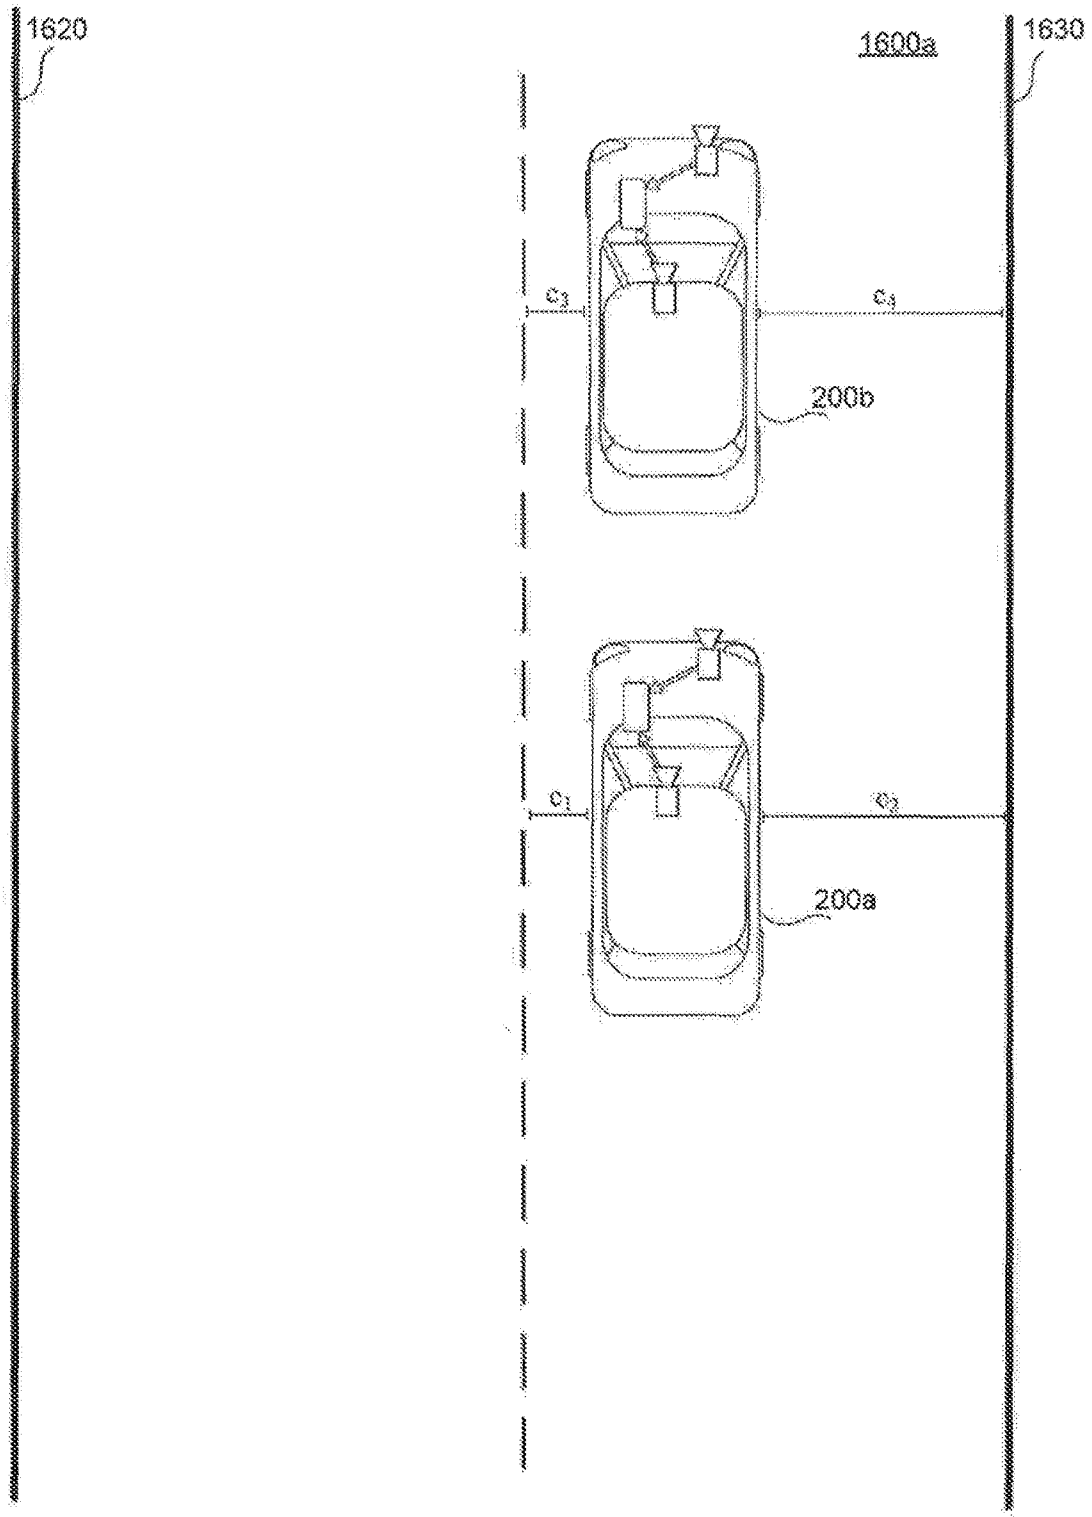


图16B

IPR2025-00943
Tesla EX1002 Page 1344

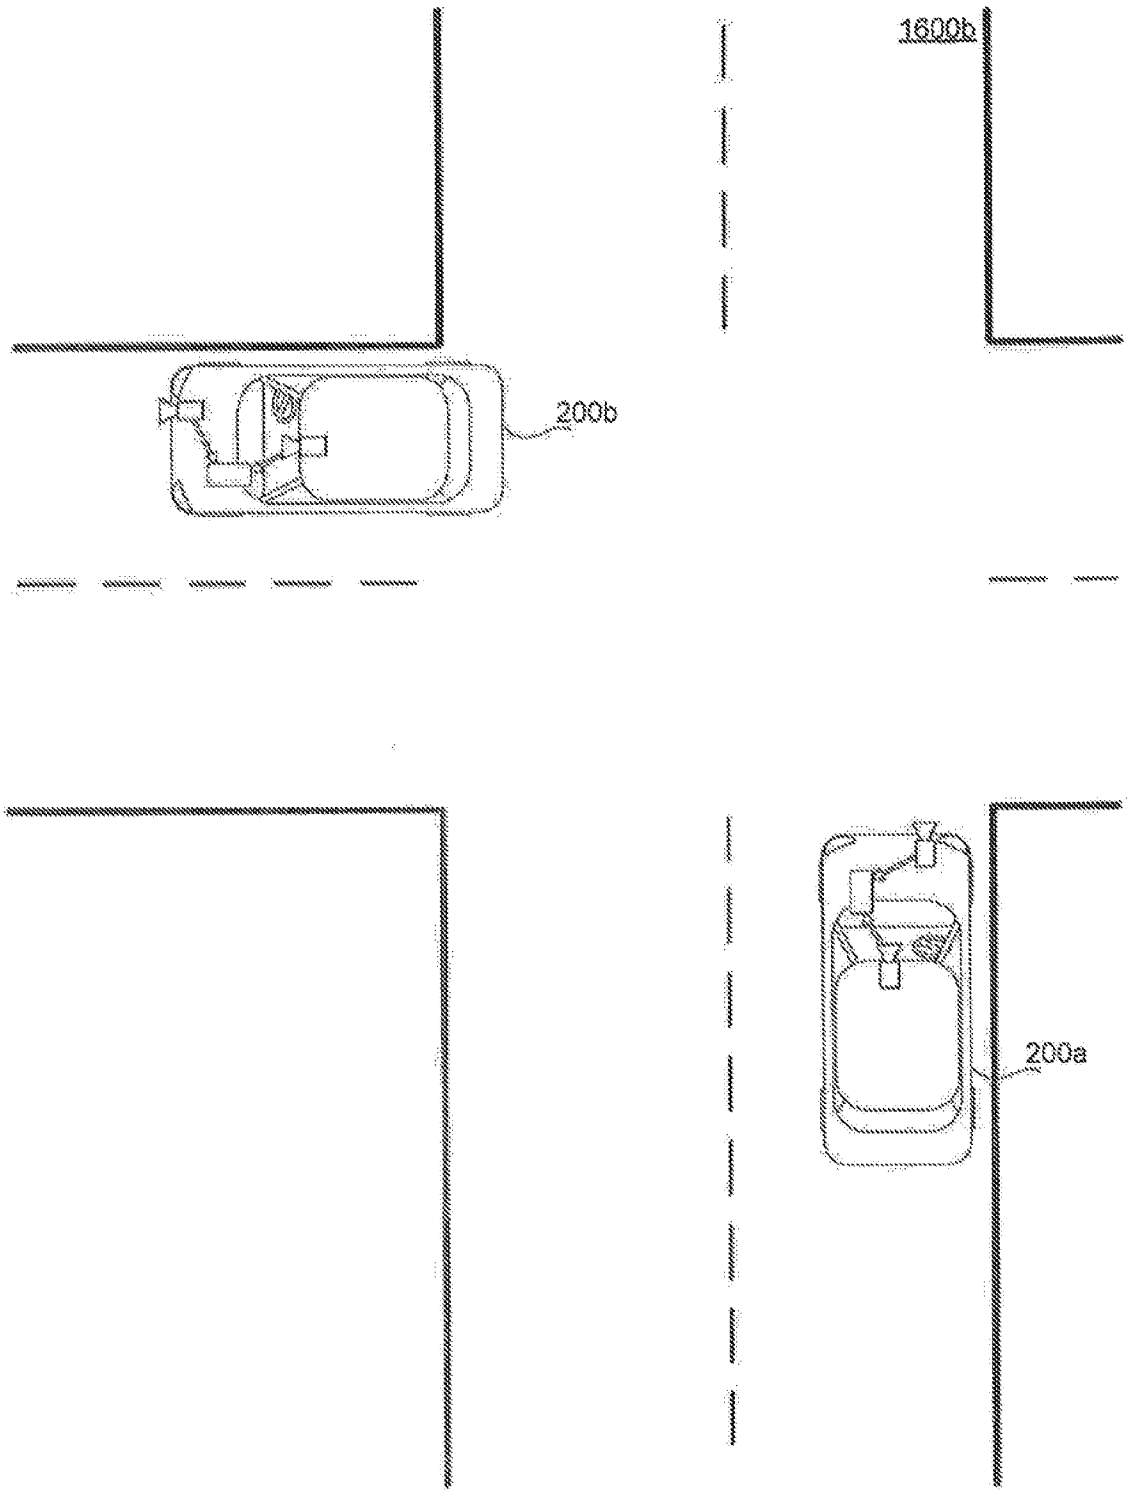


图16C

IPR2025-00943
Tesla EX1002 Page 1345

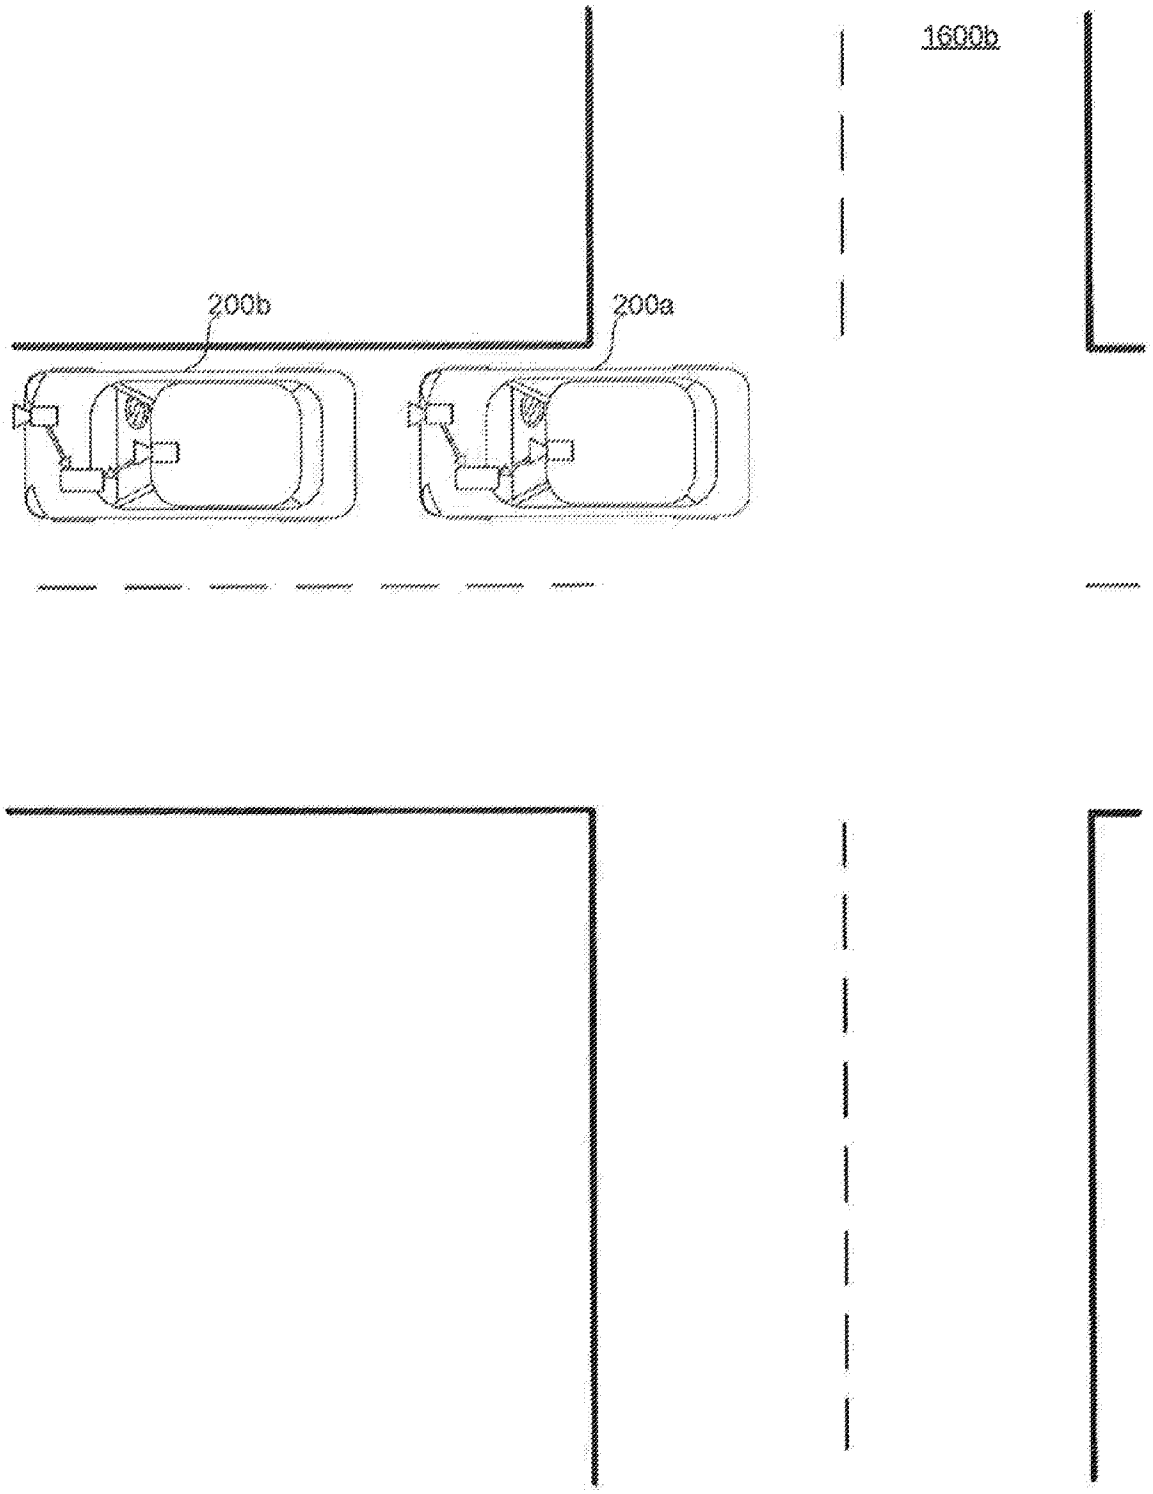


图16D

IPR2025-00943
Tesla EX1002 Page 1346

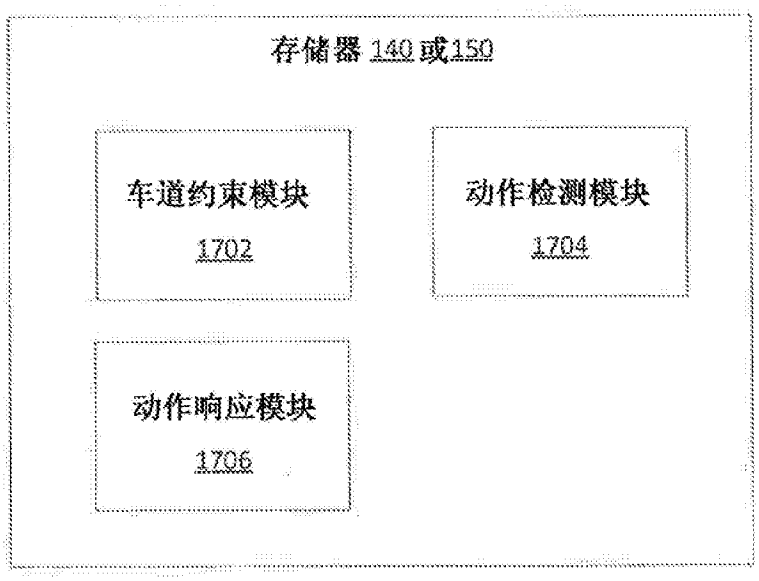


图17

IPR2025-00943
Tesla EX1002 Page 1347

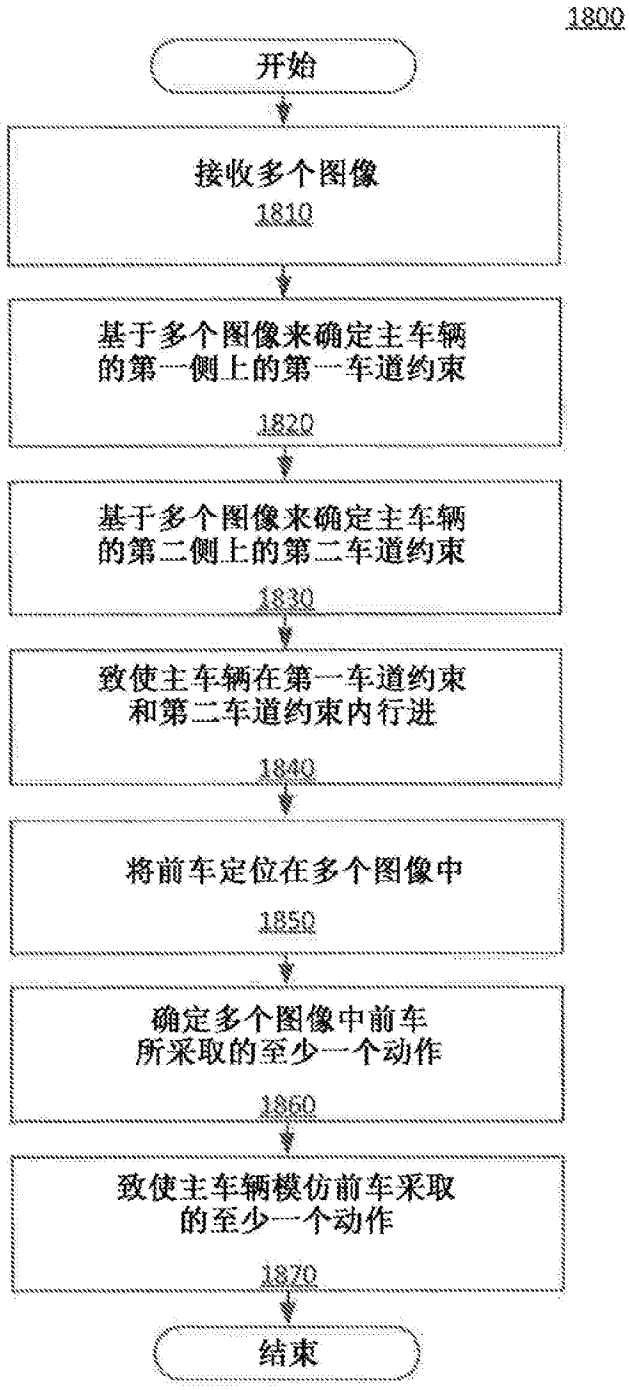


图18

IPR2025-00943
Tesla EX1002 Page 1348

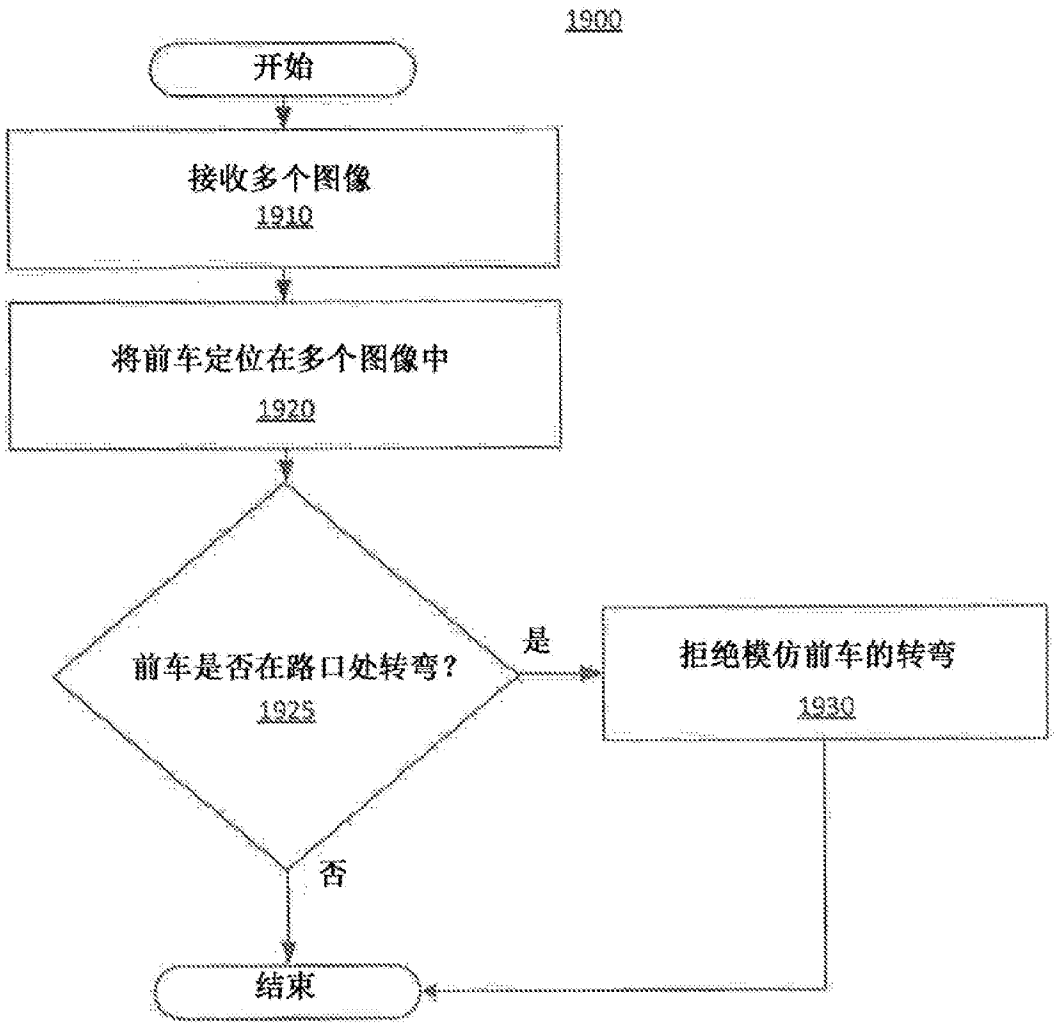


图19

IPR2025-00943
Tesla EX1002 Page 1349

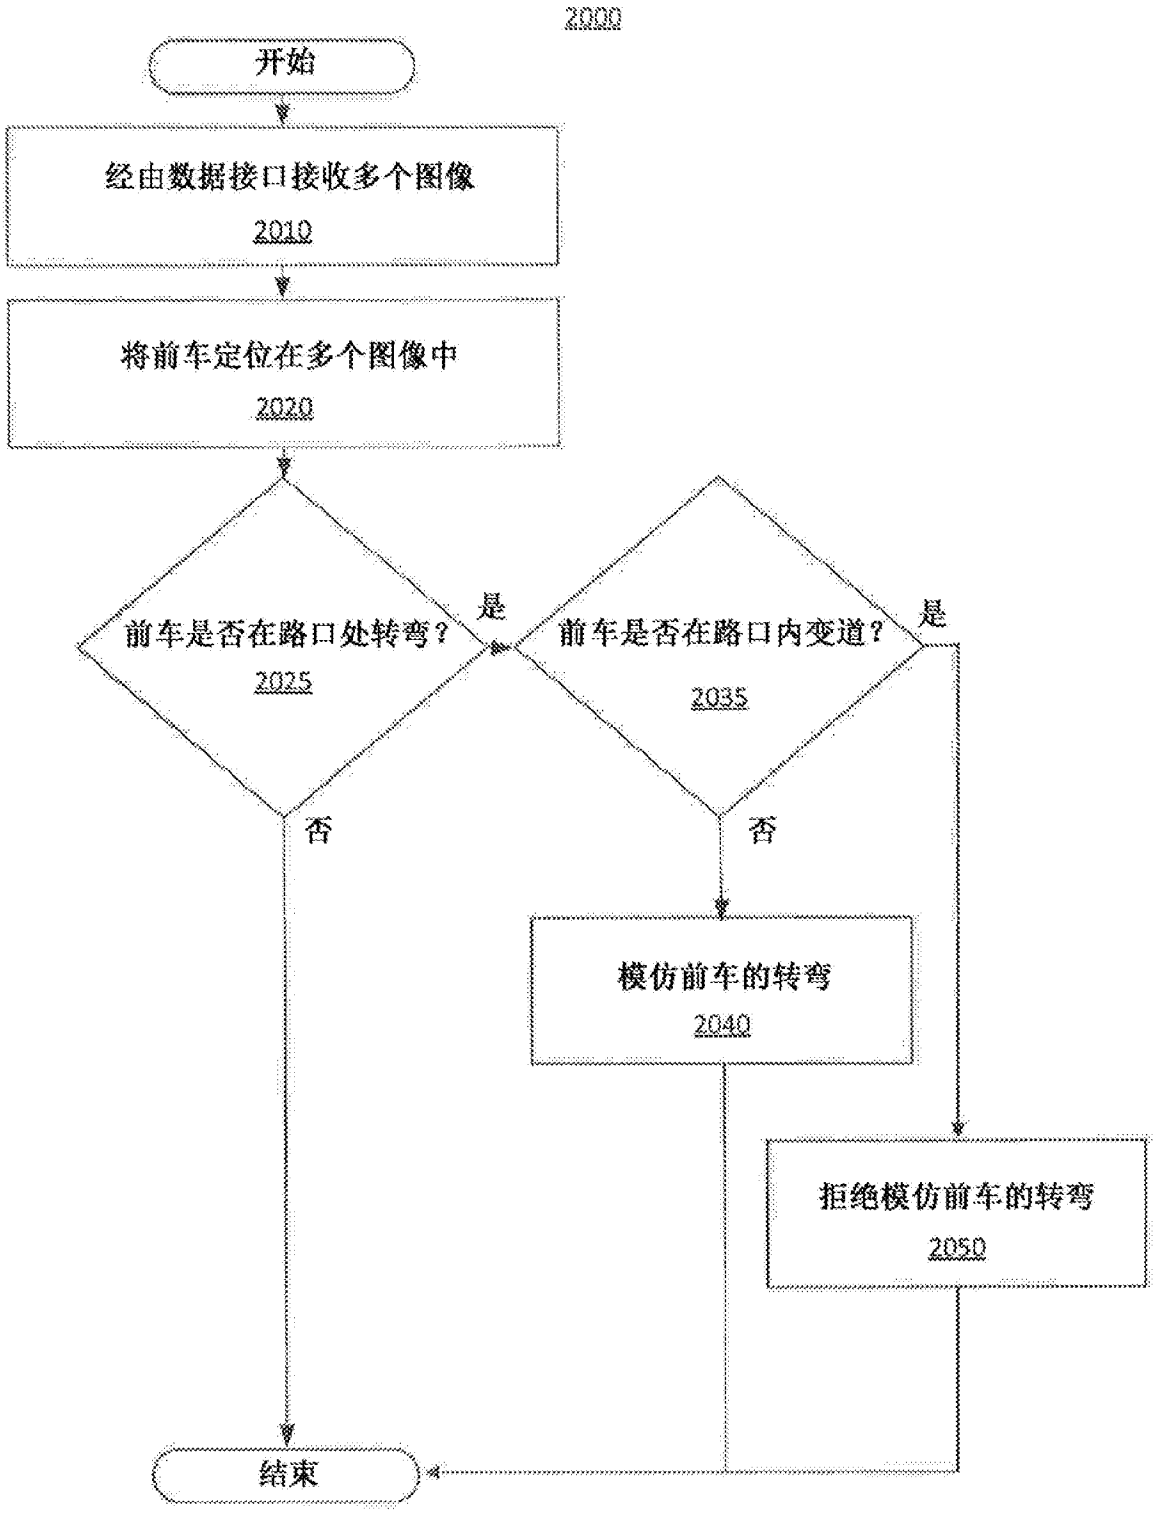


图20

IPR2025-00943
Tesla EX1002 Page 1350

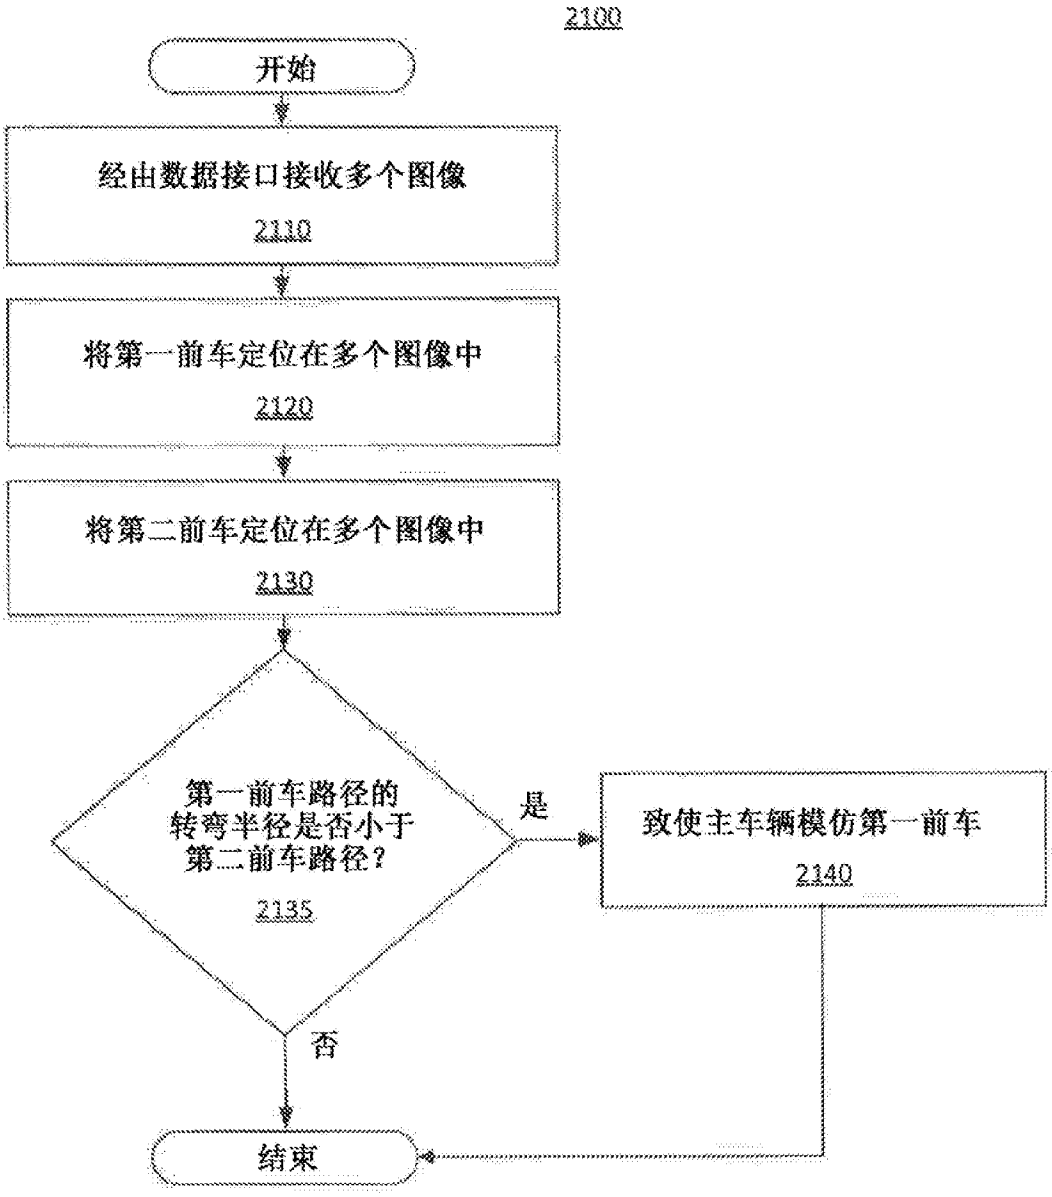


图21

IPR2025-00943
Tesla EX1002 Page 1351

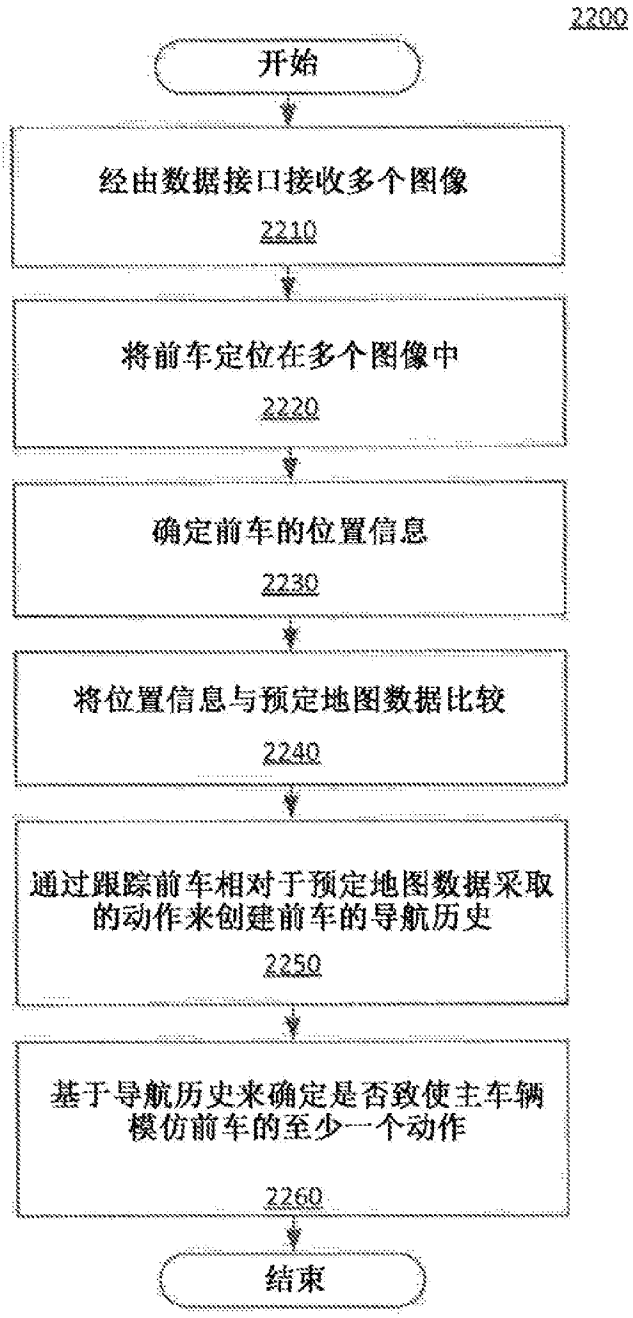


图22

IPR2025-00943
Tesla EX1002 Page 1352

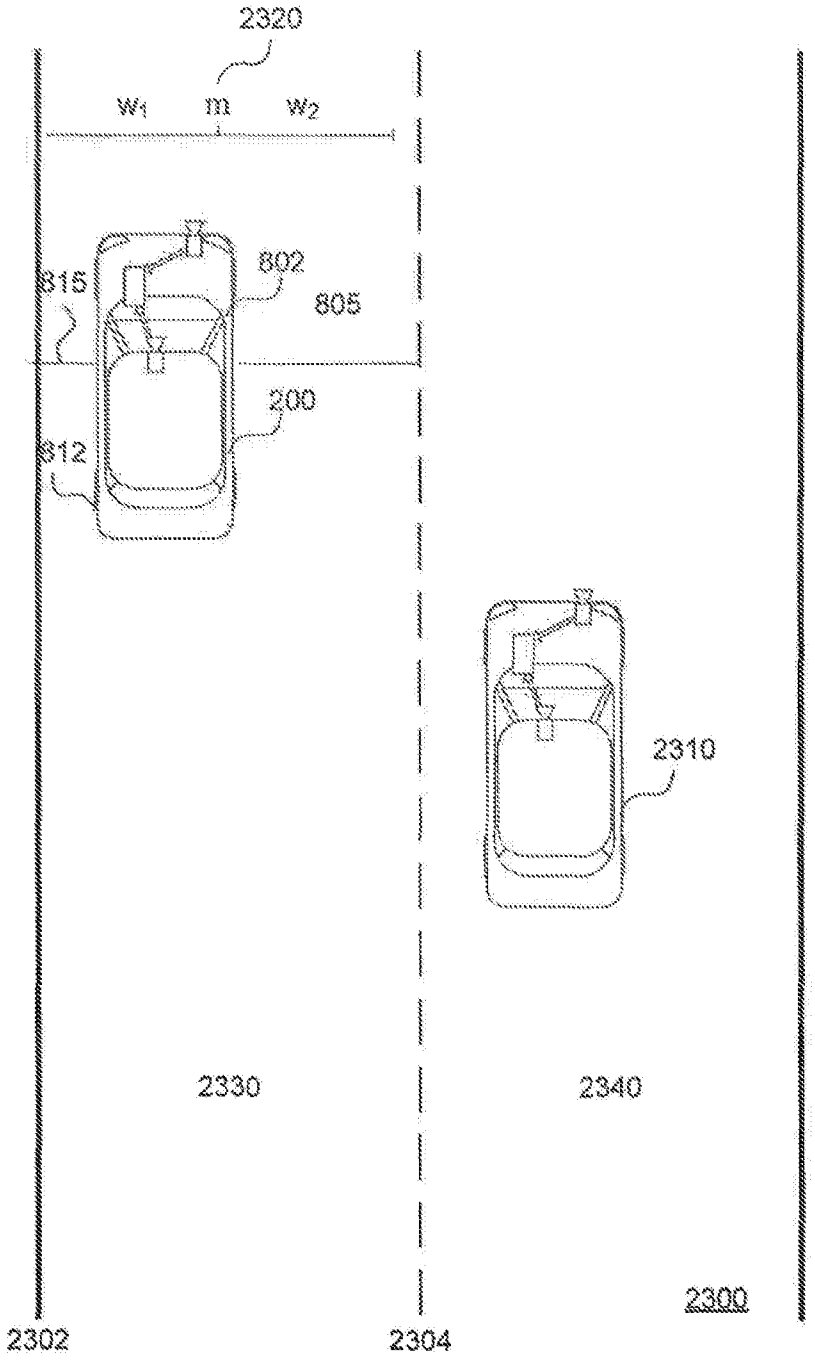


图23

IPR2025-00943
Tesla EX1002 Page 1353

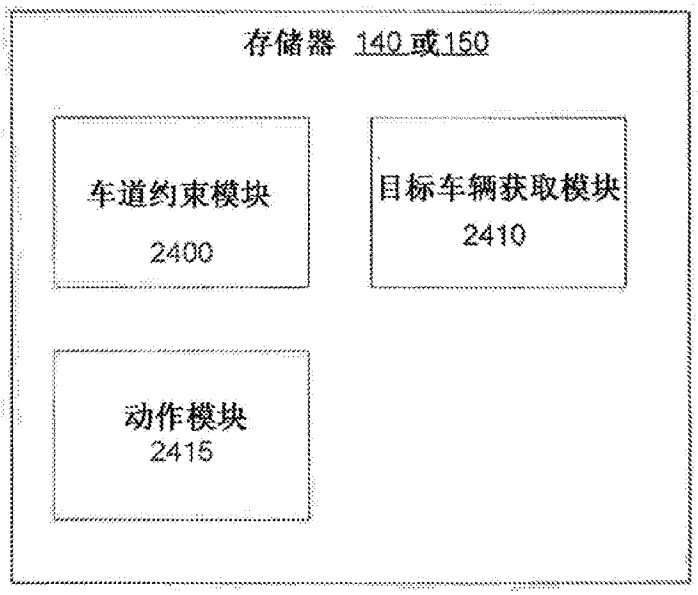


图24A

IPR2025-00943
Tesla EX1002 Page 1354

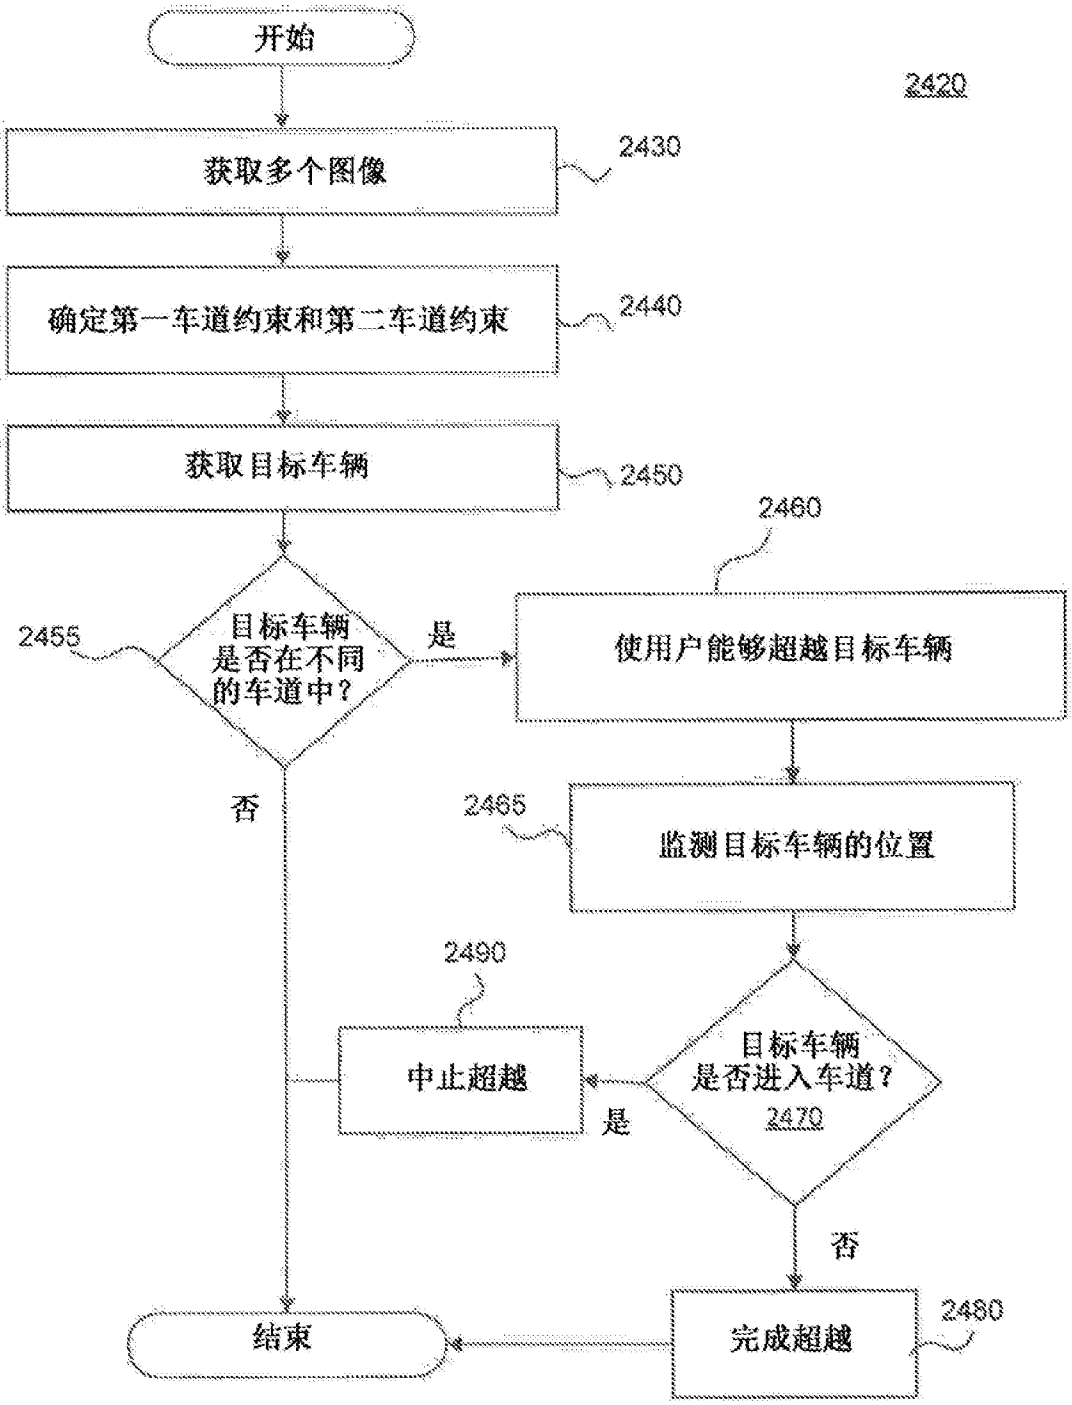


图24B

IPR2025-00943
Tesla EX1002 Page 1355

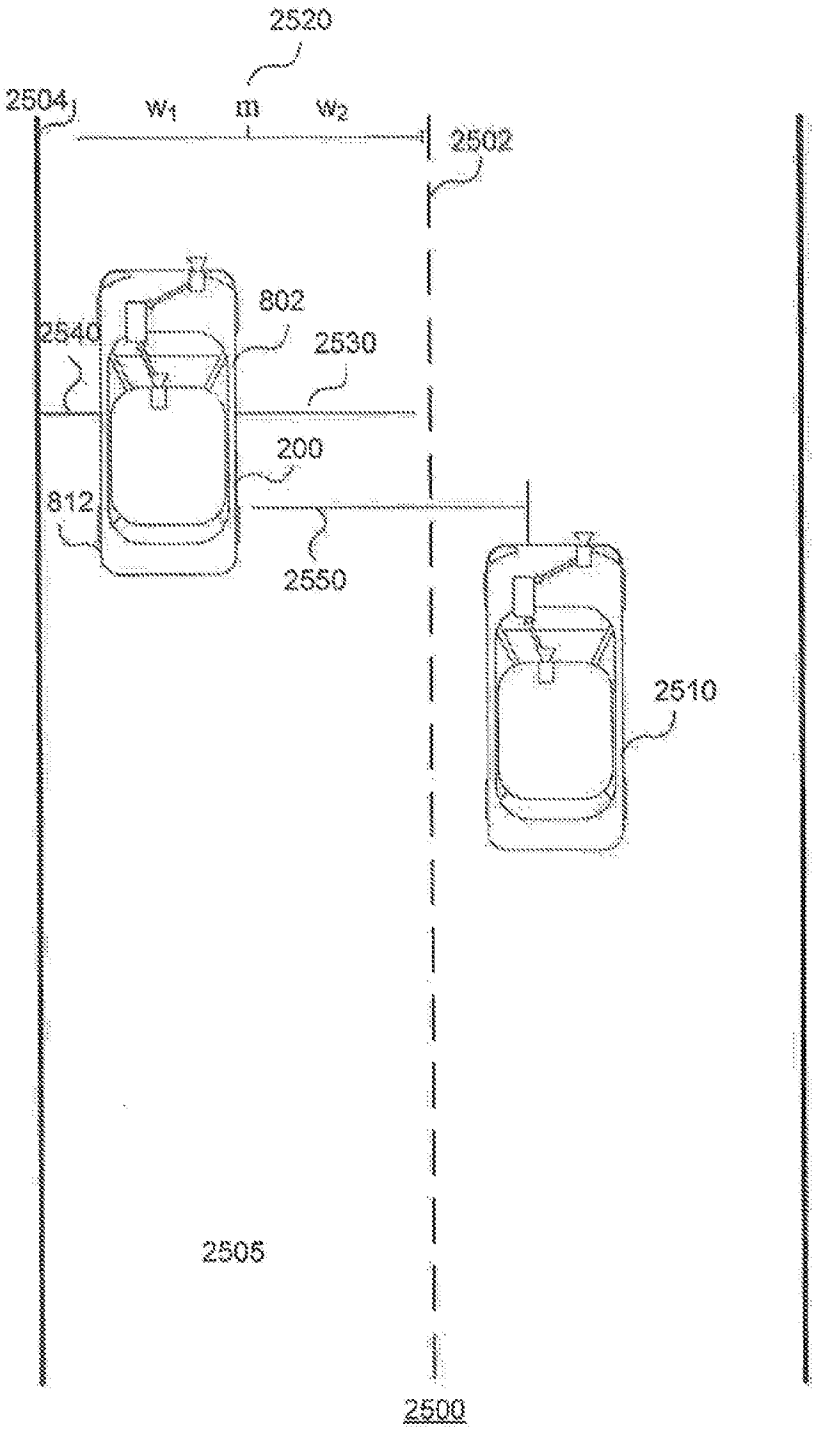


图25

IPR2025-00943
Tesla EX1002 Page 1356

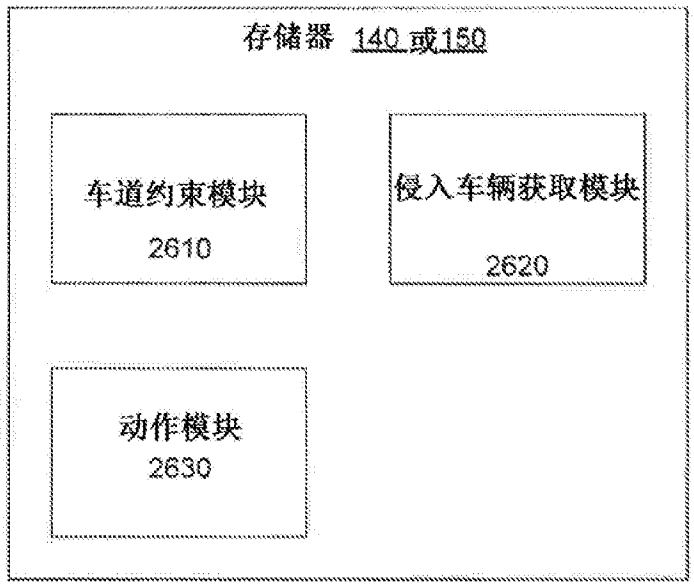


图26

IPR2025-00943
Tesla EX1002 Page 1357

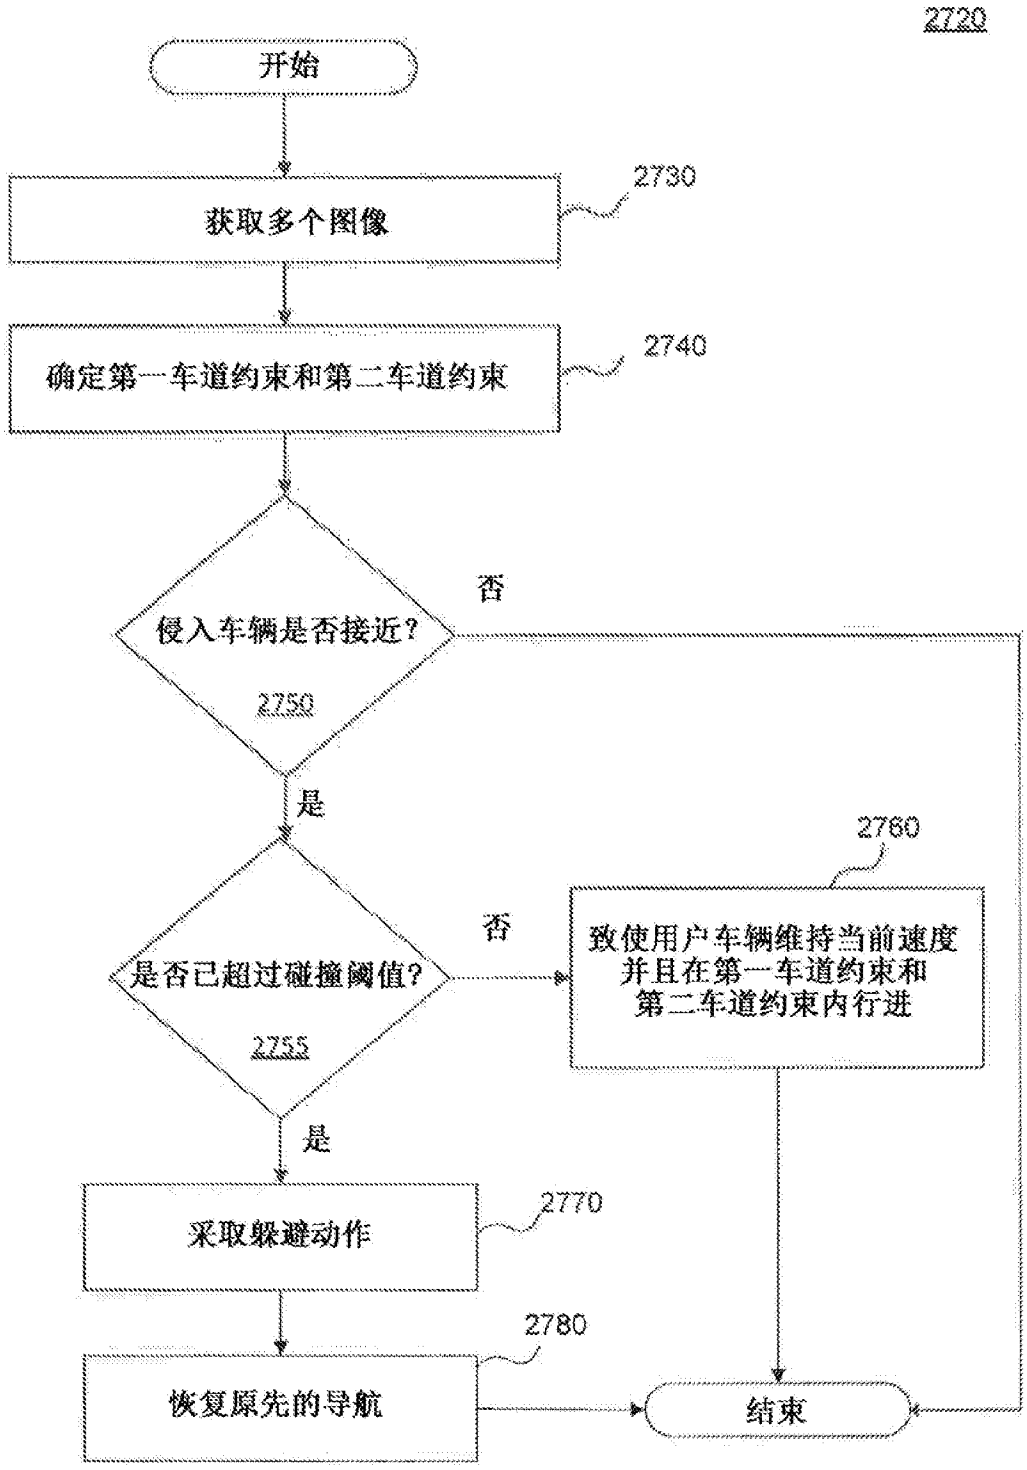


图27

IPR2025-00943
Tesla EX1002 Page 1358

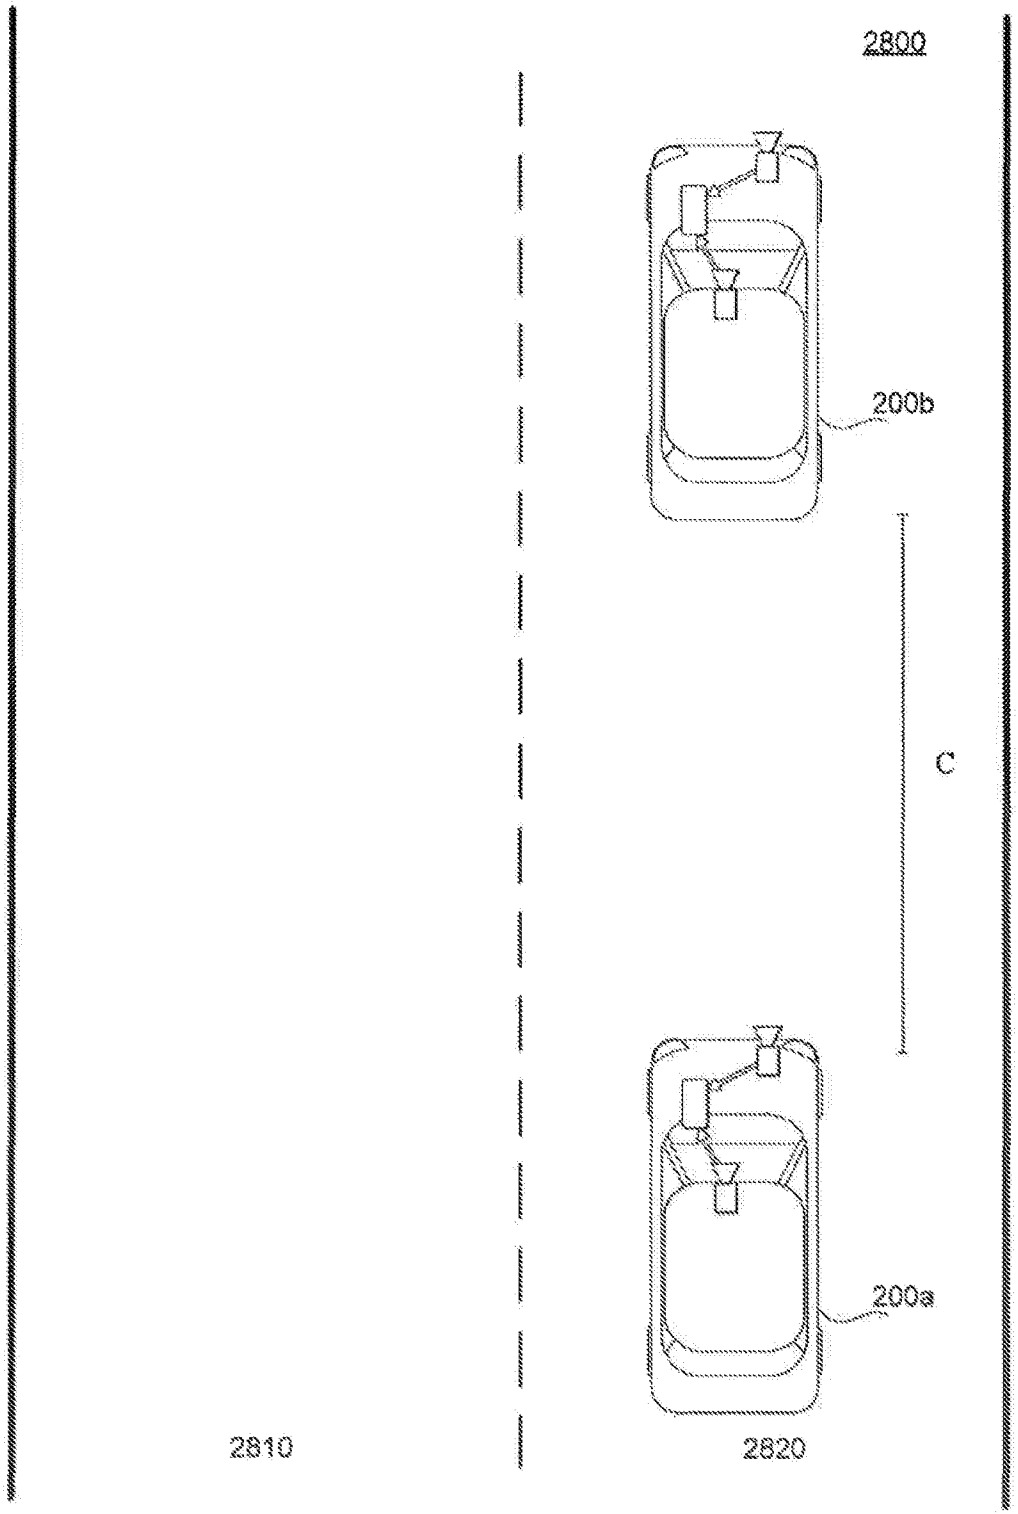


图28A

IPR2025-00943
Tesla EX1002 Page 1359

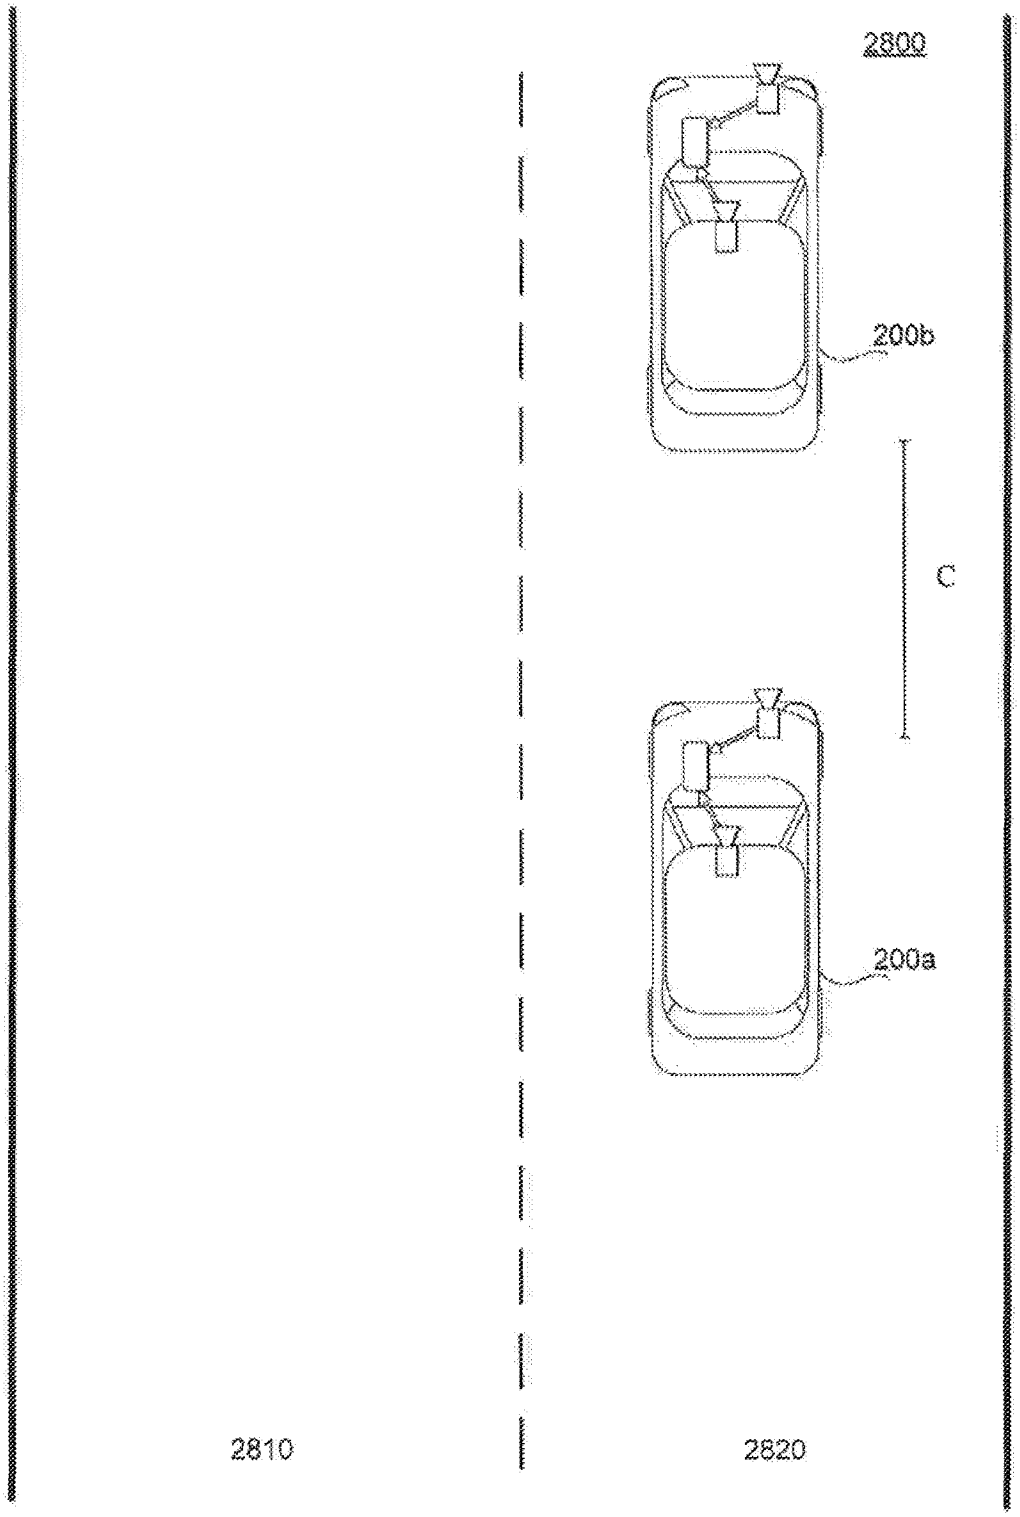


图28B

IPR2025-00943
Tesla EX1002 Page 1360

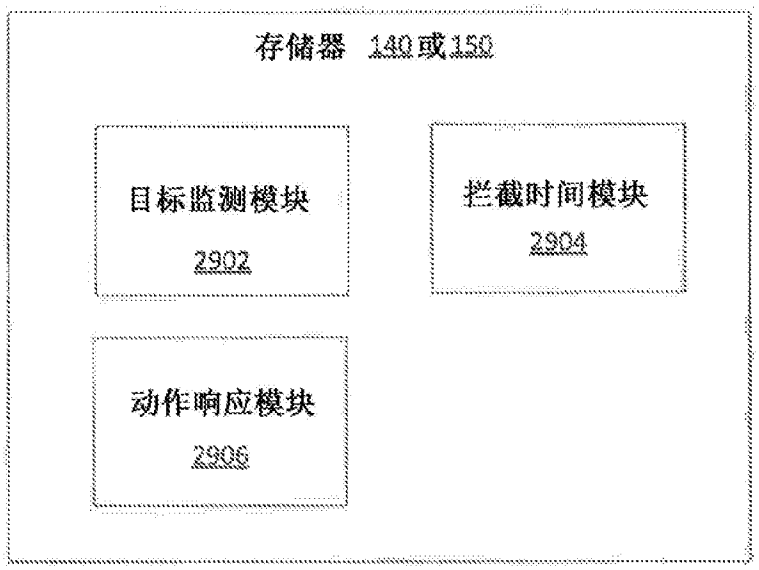


图29

IPR2025-00943
Tesla EX1002 Page 1361



图30

IPR2025-00943
Tesla EX1002 Page 1362

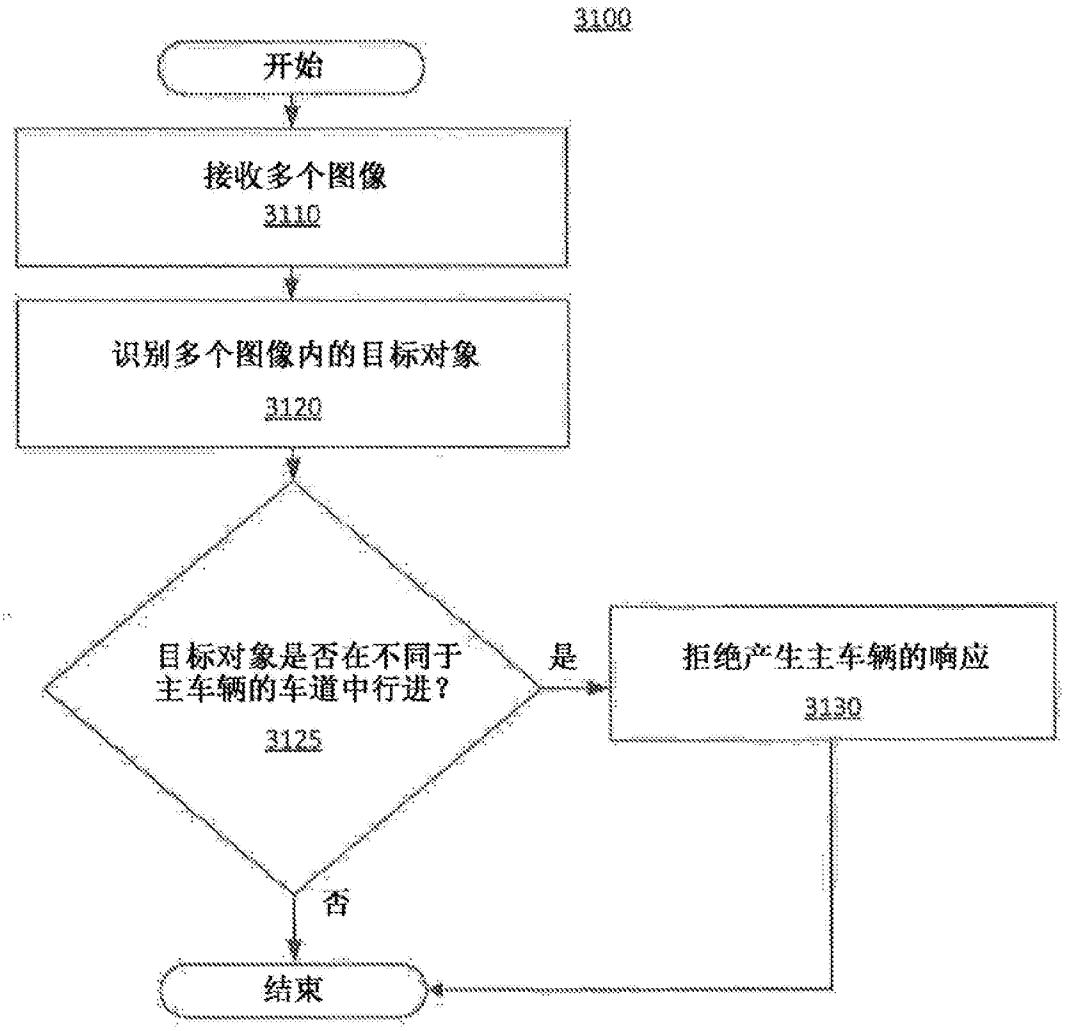


图31

IPR2025-00943
Tesla EX1002 Page 1363

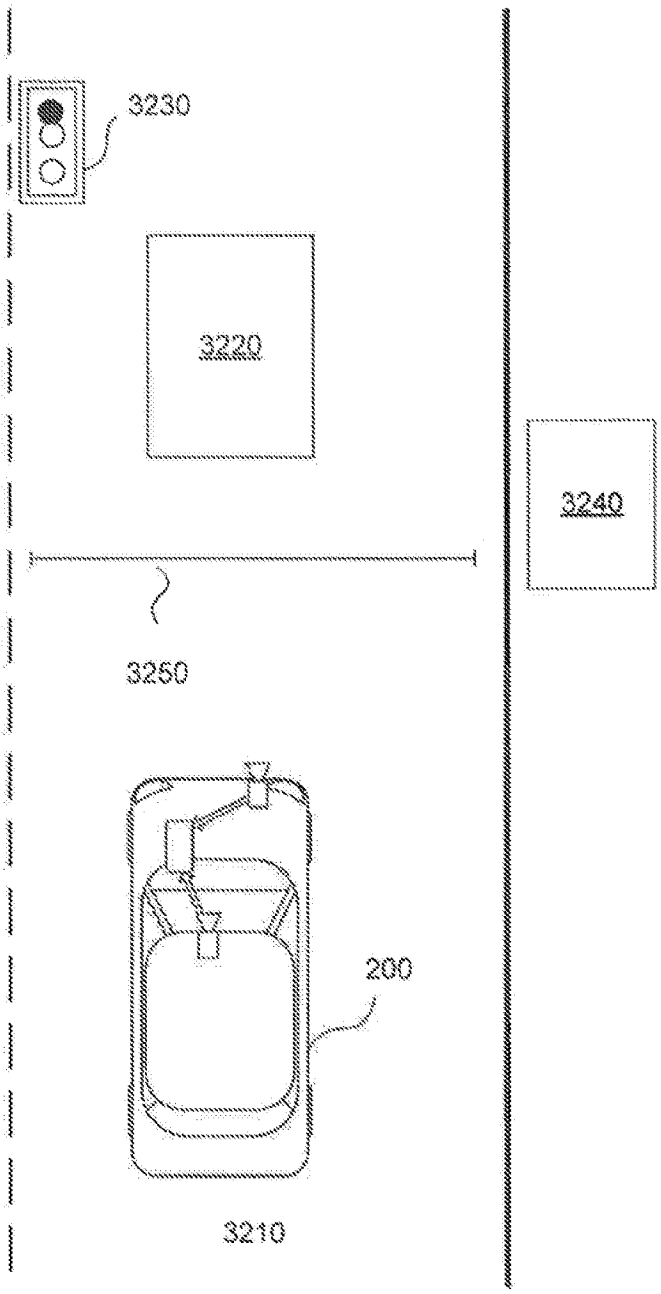


图32A

IPR2025-00943
Tesla EX1002 Page 1364

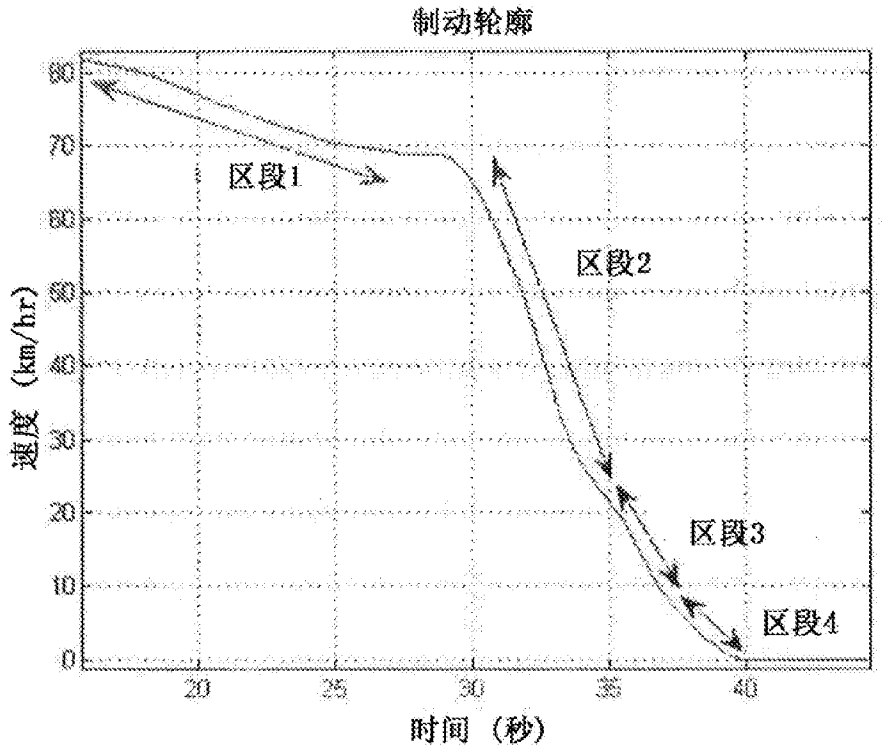


图32B



图33

IPR2025-00943
Tesla EX1002 Page 1365

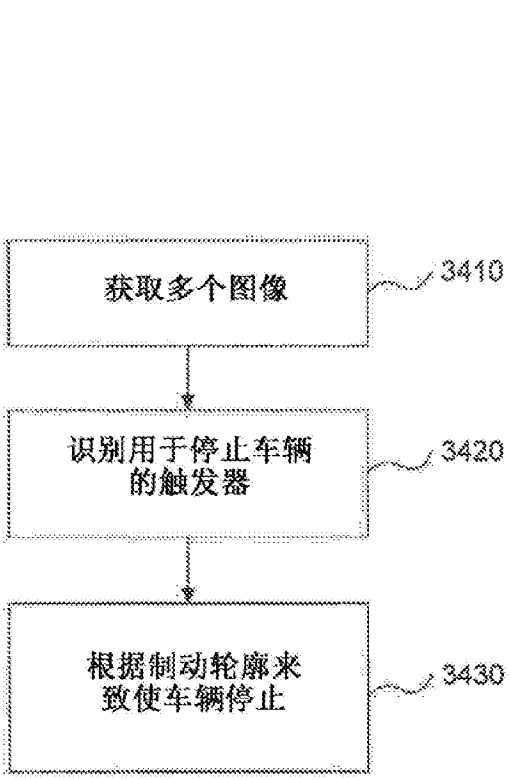


图34

IPR2025-00943
Tesla EX1002 Page 1366

Espacenet

Bibliographic data: CN201004265 (Y) — 2008-01-09

Taxi loss prevention prompt device

Inventor(s): FADIAN XU [CN] ± (XU FADIAN)

Applicant(s): FADIAN XU [CN] ± (XU FADIAN)

Classification: **- international:** ***G08B21/24; G08B7/00***
- cooperative:

Application number: CN20072006386U 20070213

Priority number(s): CN20072006386U 20070213

Abstract of CN201004265 (Y)

The utility model is a taxi property lost proof indicating device, relating to a vehicle adopting electric device, which solves the problem that the sound indication is not enough to remind the passenger. The taxi property lost proof indicating device of the utility model comprises a sound generator with a sound generating chip and a loudhailer; the sound generator is serially connected with a switch to connect the storage battery on the taxi to form a serial loop; the sound generator is also paralleily connected with an electromagnetic vibrator which is arranged on the seat of the passengers; and the switch is arranged at the opening handle in the taxi. When the switch is switched on, and the sound generator begins to work, the electromagnetic vibrator on the seat of the passenger vibrates to stimulate the passenger suddenly so as to lead the attention of the passenger into the taxi; then the passenger notices the indication of the sound generator so as to reach the effect of preventing forgetting the personal effects and baggage.

IPR2025-00943
Tesla EX1002 Page 1367

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

DESCRIPTION CN201004265Y

Taxi Lost Object Prompting Device

[0001]

Technical field

[0002]

The utility model relates to an electrical device for vehicles.

[0003]

Background technique

[0004]

When taking a taxi, it is inevitable that people will forget their belongings and luggage in the car due to the distraction of time or other reasons.

People have developed anti-lost notification devices, which use voice to remind passengers when they get off the bus to pay attention to clearing their belongings and luggage. Therefore, the existing taxi anti-lost item notification device has a voice generator with a voice generating chip and a speaker. The voice generator is connected in series with a switch in series with the battery on the car to form a series circuit. When the passenger opens the door and gets off the car, the switch is turned on and the voice generator emits a prompt voice. However, there are still some passengers who can't bother with the prompt voice. Therefore, the inventor believes that it is necessary to take certain technical measures to arouse passengers' alertness while the voice generator is working, and pay attention to the prompt voice, in order to achieve the purpose of preventing passengers from losing items.

[0005]

Utility model content

[0006]

The utility model aims to provide a prompt device for preventing lost objects in a taxi, which can arouse passengers' alertness and attention prompt voice.

[0007]

The technical scheme of the utility model is: the taxi anti-lost item prompting device has a voice generator with a voice generating chip and a speaker, the voice generator is connected in series with a switch and the battery on the car to form a series loop; the voice generator is also An electro vibrator is connected in parallel, and the electro vibrator is installed on the passenger seat.

[0008]

The switch is installed at the door opening handle in the car.

[0009]

The voice generator is installed on the car door.

[0010]

In the utility model for the taxi anti-lost notification device, an electro-induced vibrator is connected in parallel with the voice generator, and the electro-induced vibrator is installed on the passenger seat.

When the switch is turned on and the voice generator is working, the electro-vibrator on the passenger seat will vibrate, which will suddenly give the passenger a stimulus, make him alert, direct his (her) attention to the car, and he will pay attention. The prompt played by the voice generator can achieve the effect of avoiding forgetting belongings and luggage.

The switch is installed at the door opening handle in the car, and the voice generator is installed on the door. Due to the limited space in the car, passengers must keep sitting when pulling the door opening handle in the car; when the switch is turned on, the electro-vibrator vibrates The stimulation effect is the most obvious, and the voice generator is close to the door handle in the car.

The prompt effect is also the best.

[0011]

Description of the drawings

[0012]

Fig. 1 is a schematic diagram of the circuit structure of an embodiment of the present utility model of the taxi anti-lost article prompting device.

[0013]

Fig. 2 is a schematic diagram of the circuit structure of another embodiment of the lost article prompting device for a taxi of the present invention.

[0014]

Detailed ways

[0015]

One, embodiment one

[0016]

The circuit structure of an embodiment of the lost article prompting device for taxis of the present invention is shown in FIG. 1.

The taxi anti-lost notification device is composed of a switch 1, a voice generator with a voice generating chip and a speaker 2, an electromagnetic vibrator 3 and a fuse 4.

The voice generator 2 is connected in parallel with the electromagnetic vibrator 3, and then the switch 1 and the fuse 4 are connected in series to form a series circuit with the battery on the vehicle.

The switch 1 is installed at the door opening handle in the car, and the electromagnetic vibrator 3 is installed on the passenger seat near the door.

The voice generator is installed on the door.

[0017]

When the vehicle arrives at the destination and passengers want to get off, they must pull the door handle inside the vehicle.

At this time, the switch 1 is turned on and the voice generator 2 works; at the same time, the electromagnetic vibrator 3 on the passenger seat will vibrate, which will suddenly give the passenger a stimulus, make him alert, and guide his (her) attention to the car. You will notice the prompts played by the voice generator 2 to avoid forgetting your belongings and luggage.

[0018]

Second, the second embodiment

[0019]

The circuit structure of another embodiment of the lost article prompting device for taxis of the present invention is shown in FIG. 2.

The taxi anti-lost item notification device is composed of a front door switch 11, a rear door switch 12, a voice generator 20 with a voice generating chip and a speaker, three electromagnetic vibrators 30 and a fuse 40.

The voice generator 20 is connected in parallel with three electromagnetic vibrators 30 to form a prompt combination. The front door switch 11 and the rear door switch 12 are connected in parallel to form a switch combination; the prompt combination series switch combination and the fuse 40 are connected to the vehicle battery to form a series circuit.

The front door switch 11 is installed at the front door handle on the right side of the car; the rear door switch 12 is installed at the rear door handle on the right side of the car, an electromagnetic vibrator 30 is installed on the front passenger seat, and the other two electromagnetic vibrators 30 are respectively installed on the two passenger seats in the rear row.

The voice generator 20 is installed on the front door or the rear door on the right side of the vehicle.

[0020]

When the vehicle arrives at the destination and passengers want to get off, they must pull the door handle inside the vehicle.

At this time, either switch 11 or switch 12 is turned on, and the voice generator 20 can work; at the same time, the three electromagnetic vibrators 30 on the passenger seat will vibrate, and they will suddenly stimulate the passengers sitting in the corresponding seats at the same time. It arouses vigilance and directs their attention to the car, and they will notice the prompts played by the voice generator 2 to avoid forgetting their belongings and luggage.

[0021]

Obviously, although this embodiment uses one more switch and two electromagnetic vibrators than the previous embodiment, its design is more realistic and the prompt effect is better.

[0022]

The electromagnetic vibrator of the utility model of the lost object reminding device for taxis is recommended to adopt a size of about 50 mm in diameter and 5 mm in thickness. It can be directly installed on the seat cushion or backrest, or pre-fixed on the passenger's seat cover , Fixed on the passenger seat together with the seat cover.

Of course, the electromagnetic vibrator can also be replaced by a piezoelectric vibrator.

[0023]

The above are only preferred embodiments of the present utility model, but do not limit the scope of implementation of the present utility model. The equivalent changes and modifications made according to the technical scheme of the utility model and the contents of the specification shall all belong to the present utility model. The scope of the new coverage.

IPR2025-00943
Tesla EX1002 Page 1371

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

CLAIMS CN201004265Y

1.

The taxi anti-lost item notification device has a voice generator with a voice generation chip and a speaker. The voice generator is connected in series with a switch and the battery on the car to form a series circuit; it is characterized in that the voice generator is also connected in parallel with an electro-actuator Vibrator, the electro-induced vibrator is installed on the passenger seat.

2.

The taxi lost article prompting device according to claim 1, wherein the switch is installed at the door opening handle in the car.

3.

The taxi lost article prompting device according to claim 1 or 2, characterized in that: the voice generator is installed on the door of the vehicle.

[19] 中华人民共和国国家知识产权局

[51] Int. Cl.
G08B 21/24 (2006.01)
G08B 7/00 (2006.01)

[12] 实用新型专利说明书

专利号 ZL 200720006386.2

[45] 授权公告日 2008 年 1 月 9 日

[11] 授权公告号 CN 201004265Y

[22] 申请日 2007.2.13
[21] 申请号 200720006386.2
[73] 专利权人 许法典
地址 361000 福建省厦门市思明区前埔村 30 号
[72] 发明人 许法典

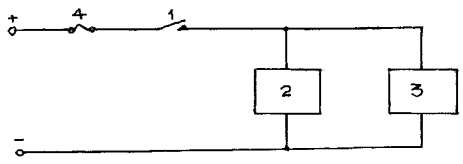
[74] 专利代理机构 厦门市首创君合专利事务所有限公司
代理人 潘国庆 张松亭
权利要求书 1 页 说明书 3 页 附图 1 页

[54] 实用新型名称

出租车防失物提示装置

[57] 摘要

本实用新型出租车防失物提示装置，涉及一种车用电器装置。解决语音提示不足以引起乘客注意的问题。该出租车防失物提示装置，有一个带语音发生芯片和扬声器的语音发生器，该语音发生器串联一个开关与车上的蓄电池连接成串联回路；该语音发生器还并联有电磁振动器，该电磁振动器安装在乘客座位上。所述的开关安装在车内的开门把手处。当开关接通，语音发生器工作的同时，乘客座位上的电磁振动器振动，会突然给乘客一个刺激，使之引起警觉，将他(她)的注意力引导到车内，就会留意到语音发生器播放的提示，达到避免遗忘随身物品和行李的效果。



IPR2025-00943
Tesla EX1002 Page 1373

1. 出租车防失物提示装置，有一个带语音发生芯片和扬声器的语音发生器，该语音发生器串联一个开关与车上的蓄电池连接成串联回路；其特征在于：该语音发生器还并联有电致振动器，该电致振动器安装在乘客座位上。

2. 根据权利要求 1 所述的出租车防失物提示装置，其特征在于：所述的开关安装在车内的开门把手处。

3. 根据权利要求 1 或 2 所述的出租车防失物提示装置，其特征在于：所述的语音发生器安装在车门上。

IPR2025-00943
Tesla EX1002 Page 1374

出租车防失物提示装置

技术领域

本实用新型涉及一种车用电器装置。

背景技术

人们在乘坐出租车时，难免因赶时间或其它原因分散精力而把随身携带的物品及行李等遗忘在车上。人们开发出防失物提示装置，在乘客下车时用声音提示他（她）们，注意清理自己的物品和行李。因此，现有的出租车防失物提示装置，都有一个带语音发生芯片和扬声器的语音发生器，该语音发生器串联一个开关与车上的蓄电池连接成串联回路。在乘客开门下车时，开关接通，语音发生器发出提示语音。但是，还有一些乘客并不能理会提示语音。从而本发明人认为：特别需要采取某种技术措施，在语音发生器工作的同时引起乘客的警觉，注意提示语音，方可达到防止乘客遗失物品的目的。

实用新型内容

本实用新型旨在提供一种出租车防失物提示装置，它可以引起乘客的警觉注意提示语音。

本实用新型的技术方案是：出租车防失物提示装置，有一个带语音发生芯片和扬声器的语音发生器，该语音发生器串联一个开关与车上的蓄电池连接成串联回路；该语音发生器还并联有电致振动器，该电致振动器安装在乘客座位上。

所述的开关安装在车内的开门把手处。

所述的语音发生器安装在车门上。

本实用新型出租车防失物提示装置，在语音发生器上并联电致振动器，该电致振动器安装在乘客座位上。当开关接通，语音发生器工作的同时，乘客座位上的电致振动器振动，会突然给乘客一个刺激，使之引起警觉，将他（她）的注意力引导到车内，就会留意到语音发生器播放的提示，达到避免遗忘随身物品和行李的效果。开关安装在车内的开门把手处，和语音发生器安装在车门上，受车内空间狭小的限制，乘客拉动车内开门把手时，必然保持坐姿；此时开关接通，电

IPR2025-00943
Tesla EX1002 Page 1375

致振动器振动给人的刺激效果最明显，语音发生器又靠近车内开门把手。提示效果也最好。

附图说明

图 1 为本实用新型出租车防失物提示装置一个实施例的电路结构示意图。

图 2 为本实用新型出租车防失物提示装置又一个实施例的电路结构示意图。

具体实施方式

一、实施例一

本实用新型出租车防失物提示装置一个实施例的电路结构，如图 1 所示。该出租车防失物提示装置，由一个开关 1、一个带语音发生芯片和扬声器的语音发生器 2、一个电磁振动器 3 和一个保险丝 4 组成。该语音发生器 2 与电磁振动器 3 并联，然后串联开关 1 和保险丝 4，与车上的蓄电池连接成串联回路。开关 1 安装在车内的开门把手处，电磁振动器 3 安装在靠近该车门的乘客座位上。语音发 2 生器安装在车门上。

当车辆到达目的地，乘客欲下车时，必须拉动车内的开门把手。此时开关 1 接通，语音发生器 2 工作；同时，乘客座位上的电磁振动器 3 振动，会突然给乘客一个刺激，使之引起警觉，将他（她）的注意力引导到车内，就会留意到语音发生器 2 播放的提示，达到避免遗忘随身物品和行李的效果。

二、实施例二

本实用新型出租车防失物提示装置又一个实施例的电路结构，如图 2 所示。该出租车防失物提示装置，由一个前门开关 11、一个后门开关 12、一个带语音发生芯片和扬声器的语音发生器 20、三个电磁振动器 30 和一个保险丝 40 组成。该语音发生器 20 与三个电磁振动器 30 并联组成提示组合，前门开关 11 与后门开关 12 并联组成开关组合；提示组合串联开关组合及保险丝 40，与车上的蓄电池连接成串联回路。前门开关 11 安装在车内右侧的前开门把手处；后门开关 12 安装在车内右侧的后开门把手处，一个电磁振动器 30 安装在前派乘客的座位上，另外两个电磁振动器 30 分别安装在后排的两个乘客座席位置上。语音发生器 20 安装在车内右侧的前车门或后车门上。

IPR2025-00943
Tesla EX1002 Page 1376

当车辆到达目的地，乘客欲下车时，必须拉动车内的开门把手。此时开关11或开关12中的任何一个接通，语音发生器20都可工作；同时，乘客座位上的三个电磁振动器30振动，会同时突然给坐在对应坐席的乘客一个刺激，使之引起警觉，将他（她）们的注意力引导到车内，就会留意到语音发生器2播放的提示，达到避免遗忘随身物品和行李的效果。

显然本实施例虽比前一实施例多用了一个开关和两个电磁振动器，但它的设计更合乎实际，提示效果更好。

本实用新型出租车防失物提示装置的电磁振动器，建议采用直径约50毫米、厚度约5毫米的尺寸，它可以直接安装在座位的坐垫或靠背上，也可以预先固定在乘客的座位套上，随座位套一起固定在乘客的座位上。当然，电磁振动器也可使用压电效应的振动器来替代。

以上所述，仅为本实用新型较佳实施例，然而并不以此限定本实用新型实施的范围，依本实用新型的技术方案及说明书内容所作的等效变化与修饰，皆应属于本实用新型涵盖的范围。

IPR2025-00943
Tesla EX1002 Page 1377

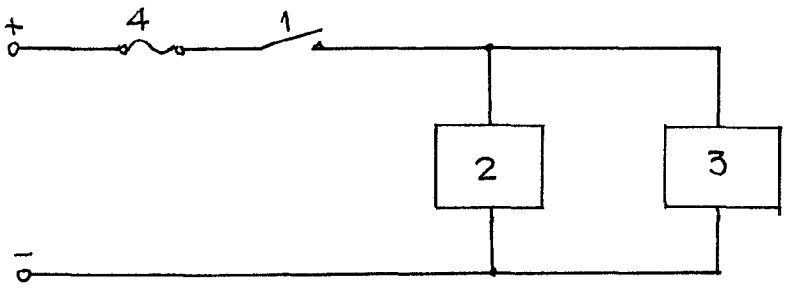


图1

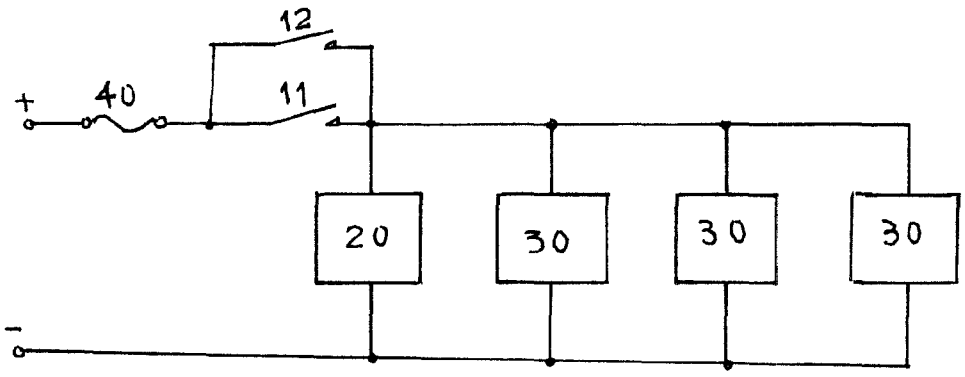


图2

IPR2025-00943
Tesla EX1002 Page 1378

Espacenet

Bibliographic data: CN202012052 (U) — 2011-10-19

Novel road gate

Inventor(s): MENGSHENG MA; LONG XU ± (MA MENGSHENG, ; XU LONG)

Applicant(s): MENGSHENG MA ± (MA MENGSHENG)

Classification: **- international:** ***E01F13/04***
- cooperative:

Application number: CN20112052465U 20110302

Priority number(s): CN20112052465U 20110302

Abstract of CN202012052 (U)

The utility model discloses a novel road gate, comprising a bracket and a rotary pole. One end of the rotary pole is arranged on the bracket, and the rotary pole is capable of rotating around the bracket. Disposed on the rotary pole, an information box comprises a box frame, battens and information bodies. The box frame is made of hollow aluminum pipes. The information bodies are pressed and arranged on front and rear side surfaces of the box frame via the battens. The information box is added on a conventional road gate, so that the road gate not only has the functions of opening and closing roads, and also has different functions of prompting advertisement information, prompting parking information of a parking area, etc. In addition, the information box is light, safe and reliable, and the information bodies on the information box are convenient to replace and use.

IPR2025-00943
Tesla EX1002 Page 1379

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

DESCRIPTION CN202012052U

New type of barrier

[0001]

Technical field

[0002]

The utility model relates to the technical field of barrier gates, in particular to a new type of barrier gates.

[0003]

Background technique

[0004]

Barrier gates are special access management equipment used to restrict motor vehicles on roads. They are now widely used in highway toll stations, parking lots, communities, and the entrances of enterprises and institutions to manage vehicle access.

At present, there are two main types of barriers, one is manual and the other is electric. The barrier is mainly composed of a bracket and a rotating support rod. One end of the rotating support rod is installed on the support, and the rotating support rod rotates around the support to realize the closing and opening of the road.

[0005]

At present, most of the barriers on the market have their main function to close and open roads, and their functions are relatively single; there are also a small number of barriers equipped with information boards on the rotating poles, which provide advertisements and information in the parking lot. Parking information, etc., but the information boards on these barriers are generally made of solid plates or plastics, making the information boards heavier, prolonged use of On the rotating support rod, the information posted on the information board is generally fixed on the information board, and there are many inconveniences in replacing the information.

[0006]

Summary of the invention

[0007]

In order to overcome the shortcomings of the prior art, the utility model provides a new type of barrier gate, which solves the shortcomings of the existing barrier gates of single function, hidden safety hazards, and inconvenient information replacement.

[0008]

In order to achieve the above objectives, the technical solutions adopted by the present utility model to solve its technical problems are:

[0009]

A new type of barrier, including a bracket and a rotating pole. One end of the rotating pole is installed on the bracket. The rotating pole can rotate around the bracket. The rotating pole is equipped with an information box, and the rotating pole crosses both ends of the information box. So that the information box is fixed on the rotating support rod.

[0010]

An information box is installed on the rotating support rod.

[0011]

The information box is composed of a box frame body, a layering bar, and a message body. The box frame body is made of a hollow aluminum tube. The message body is installed on the front and back sides of the box frame by pressing the bead, and it is very convenient to unload the bead. Change different information bodies locally.

[0012]

The display content and display form of the information body can be different, so as to achieve different functional effects.

[0013]

The beneficial effects of the utility model are:

[0014]

1、An information box is added to the rotating poles of the existing barrier gates. The content and form of the information displayed on the information boxes can be different, so that the barriers not only have the function of road closure and opening, but also have advertising information prompts and parking lot parking. Different functions such as information prompt;

[0015]

2、The box body of the information box is made of hollow aluminum tube, which makes the information box lighter and lighter, so that it is not easy to damage the bracket and is safe and reliable to use;

[0016]

3、The information body of the information box is installed on the front side and the back side of the box frame by pressing the bead, and by disassembling the bead, different information bodies can be easily replaced, which is convenient to use.

[0017]

Description of the drawings

[0018]

Figure 1 is a schematic diagram of the structure of the utility model.

[0019]

Figure 2 is a cross-sectional view of the information box of the present invention.

[0020]

Detailed ways

[0021]

In the following, the present utility model will be further explained in conjunction with specific embodiments. It should be understood that these embodiments are only used to illustrate the present utility model and not to limit the scope of the present utility model. In addition, it should be understood that after reading the teachings of the present utility model, those skilled in the art can make various changes or modifications to the present utility model, and these equivalent forms also fall within the scope defined by the appended claims of the application.

[0022]

Referring to Figures 1 and 2, the present invention includes a bracket 30 and a rotating support rod 10. One end of the rotating support rod 10 is mounted on the support 30, the rotating support rod 10 can rotate around the support 30, and an information box 20 is installed on the rotating support rod 10 , The rotating support rod 10 crosses both ends of the information box 20 so that the information box is fixed on the rotating support rod.

The information box is composed of a box body 21, a bead 222, and a message body 23. The box body 21 is made of a hollow aluminum tube. With the unloading and installation, different information bodies 23 can be easily replaced. The display content of the information body 23 can be different, it can be advertising information, parking information in a parking lot, or other information. The display form of the information body 23 can be different, it can be advertising paper, it can be an exhibition board, or other information. form.

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

CLAIMS CN202012052U

1.

A new type of barrier includes a bracket and a rotating support rod. One end of the rotating support rod is installed on the support, and is characterized in that an information box is installed on the rotating support rod.

2.

A new type of barrier gate according to claim 1, characterized in that: an information box is installed on the rotating support rod.

3.

A new type of barrier gate according to claim 1, characterized in that: the information box is composed of a box frame body, a layering bar, and an information body. The box frame body is made of a hollow aluminum tube, and the information body is installed on the box frame by pressing the layering bar. The front and back sides of the body.

(19)中华人民共和国国家知识产权局

(12)实用新型专利

(10)授权公告号 CN 202012052 U
(45)授权公告日 2011.10.19

(21)申请号 201120052465.3

(22)申请日 2011.03.02

(73)专利权人 马萌声
地址 430000 湖北省武汉市江汉区台北一路天下国际公馆1单元507

(72)发明人 马萌声 徐龙

(51)Int.Cl.
E01F 13/04(2006.01)

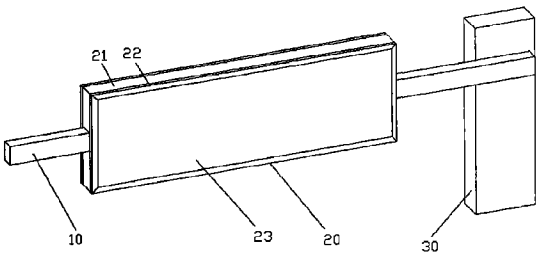
权利要求书 1 页 说明书 2 页 附图 1 页

(54)实用新型名称

新型道闸

(57)摘要

本实用新型公开了一种新型道闸，包含支架和旋转支杆，旋转支杆一端安装于支架上，旋转支杆可以绕支架旋转，旋转支杆上安装有信息箱，信息箱由箱框体、压条、信息体组成，箱框体由空心铝管制成，信息体通过压条压制安装在箱框体的前侧面和后侧面。本实用新型在现有的道闸上加设有信息箱，使得道闸不仅具有道路封闭与开放的功能，还具有广告信息提示、停车场停车信息提示等不同功能，且信息箱的重量轻，安全可靠，信息箱上的信息体便于更换，使用方便。



CN 202012052 U

IPR2025-00943
Tesla EX1002 Page 1384

1. 一种新型道闸，包含支架和旋转支杆，旋转支杆一端安装于支架上，其特征在于：旋转支杆上安装有信息箱。

2. 根据权利要求 1 所述的一种新型道闸，其特征在于：所述旋转支杆上安装有一个信息箱。

3. 根据权利要求 1 所述的一种新型道闸，其特征在于：所述信息箱由箱框体、压条、信息体组成，箱框体由空心铝管制成，信息体通过压条压制安装在箱框体的前侧面和后侧面。

IPR2025-00943
Tesla EX1002 Page 1385

新型道闸

技术领域

[0001] 本实用新型涉及道闸技术领域，特别是涉及一种新型道闸。

背景技术

[0002] 道闸是专门用于道路上限制机动车行驶的通道出入口管理设备，现广泛应用于公路收费站、停车场、小区、企事业单位门口，来管理车辆的出入。目前，道闸主要有两种，一种是手动的，另一种是电动的。道闸主要由支架和旋转支杆组成，旋转支杆一端安装于支架上，旋转支杆绕支架旋转，可以实现道路的封闭与开放。

[0003] 目前市场上的大部分道闸，主要功能是实现道路的封闭与开放，功能较为单一；也有小部分的道闸的旋转支杆上加装了信息板，信息板上提供广告、停车场内的停车信息等，但这些道闸上的信息板，一般都有实心板材或者塑料做成，使得信息板的重量较重，长时间使用易损坏支架且存在一定的安全隐患，同时，信息板固定在旋转支杆上，信息板上登载的信息一般也固定在信息板上，更换信息存在诸多不便。

发明内容

[0004] 为克服现有技术的不足，本实用新型提供一种新型道闸，解决现有道闸功能单一、存在安全隐患、信息更换不便的不足。

[0005] 为达到上述目的，本实用新型解决其技术问题所采用的技术方案是：

[0006] 一种新型道闸，包含支架和旋转支杆，旋转支杆一端安装于支架上，旋转支杆可以绕支架旋转，旋转支杆上安装有信息箱，旋转支杆横穿过信息箱的两端从而使得信息箱固定在旋转支杆上。

[0007] 所述旋转支杆上安装有一个信息箱。

[0008] 所述信息箱由箱框体、压条、信息体组成，箱框体由空心铝管制成，信息体通过压条压制安装在箱框体的前侧面和后侧面，通过对压条的卸装，可以很方便地更换不同的信息体。

[0009] 所述信息体显示内容和显示形态可以不同，从而达到不同的功能效果。

[0010] 本实用新型的有益效果是：

[0011] 1、在现有的道闸的旋转支杆上加设有信息箱，信息箱上显示的信息内容和形态可以不同，使得道闸不仅具有道路封闭与开放的功能，还具有广告信息提示、停车场停车信息提示等不同功能；

[0012] 2、信息箱的箱框体由空心铝管制成，使得信息箱重量轻，较为轻便，从而不易损坏支架且使用安全可靠；

[0013] 3、信息箱的信息体通过压条压制安装在箱框体的前侧面和后侧面，通过对压条的卸装，可以很方便地更换不同的信息体，使用方便。

附图说明

IPR2025-00943
Tesla EX1002 Page 1386

[0014] 图 1 是本实用新型结构示意图。

[0015] 图 2 是本实用新型信息箱的断面图。

具体实施方式

[0016] 下面结合具体实施例，进一步阐述本实用新型。应理解，这些实施例仅用于说明本实用新型而不用于限制本实用新型的范围。此外应理解，在阅读了本实用新型讲授的内容之后，本领域技术人员可以对本实用新型作各种改动或修改，这些等价形式同样落在申请所附权利要求书所限定的范围。

[0017] 参见图 1、图 2，本实用新型包含支架 30 和旋转支杆 10，旋转支杆 10 一端安装于支架 30 上，旋转支杆 10 可以绕支架 30 旋转，旋转支杆 10 上安装有信息箱 20，旋转支杆 10 横穿过信息箱 20 的两端从而使得信息箱固定在旋转支杆上。信息箱由箱框体 21、压条 222、信息体 23 组成，箱框体 21 由空心铝管制成，信息体 23 通过压条 22 压制安装在箱框体 21 的前侧面和后侧面，通过对压条 22 的卸装，可以很方便地更换不同的信息体 23。信息体 23 显示内容可以不同，可以是广告信息，也可以是停车场的停车信息，亦或是其他信息，信息体 23 显示形态可以不同，可以是广告纸张，也可以是展板，亦或是其他形态。

IPR2025-00943
Tesla EX1002 Page 1387

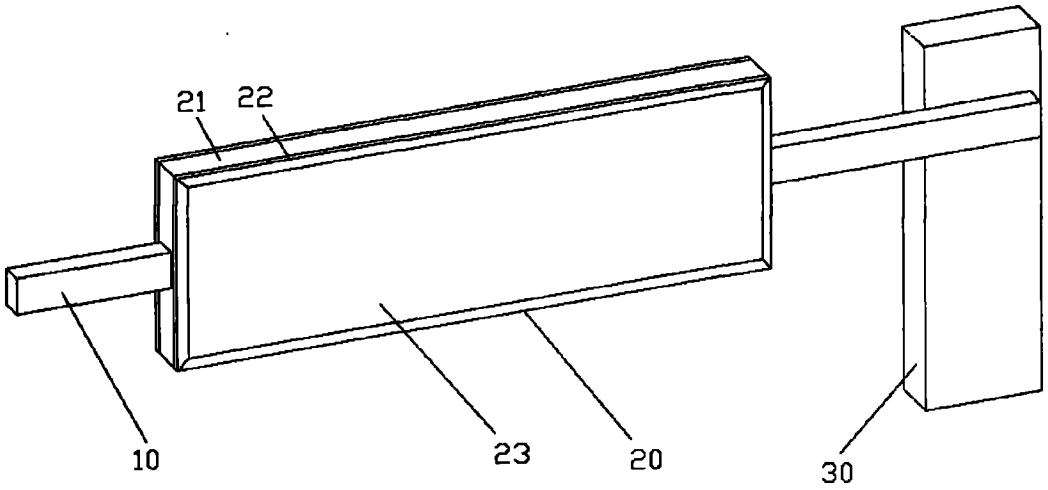


图 1

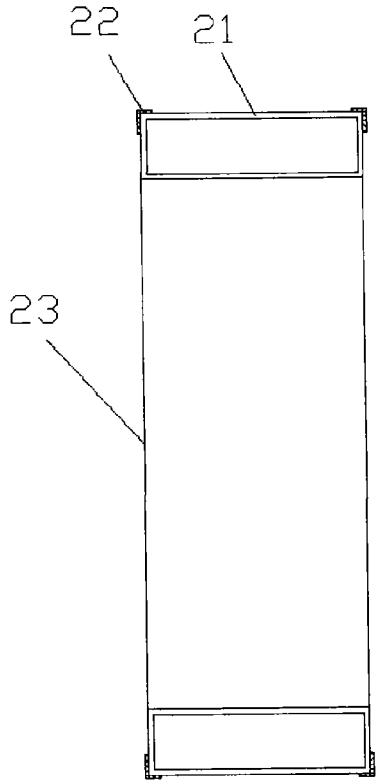


图 2

IPR2025-00943
Tesla EX1002 Page 1388

Espacenet

Bibliographic data: CN202038228 (U) — 2011-11-16

Automatic luggage reminding device for taxi trunk

Inventor(s): GUOBING LIU; YOUQUN ZHAO ± (LIU GUOBING, ; ZHAO YOUQUN)

Applicant(s): NANJING INST TECHNOLOGY ± (NANJING INSTITUTE OF TECHNOLOGY)

Classification: **- international:** ***B60Q5/00***
- cooperative:

Application number: CN201120113686U 20110418

Priority number(s): CN201120113686U 20110418

Abstract of CN202038228 (U)

The utility model relates to an automatic luggage reminding device for a taxi trunk, which comprises a power circuit, wherein, the output end of the power circuit is respectively connected with a trigger circuit, a controller and a voice circuit, the trigger circuit is composed of a first switch linked with the trunk and a second switch linked with a taximeter, the first switch is in parallel connection with the second switch, and the controller is provided with a first monitor input end connected with the first switch, a second monitor input end connected with the second switch, and an output end connected with the voice circuit. The automatic luggage reminding device can automatically send a reminding message to passengers only when necessary, without repeatedly playing voice prompt for no reason.

IPR2025-00943
Tesla EX1002 Page 1389

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

DESCRIPTION CN202038228U

Automatic reminder for luggage in the trunk of taxi

[0001]

Technical field

[0002]

The utility model relates to a special reminder device for taxis, in particular to an automatic reminder device for the trunk luggage of a taxi.

[0003]

Background technique

[0004]

According to the applicant's understanding, passengers generally put their luggage in the trunk of the car when taking a taxi, and it is easy to forget their luggage when getting off the bus.

Existing reminder devices for taxis generally turn on the voice circuit by operating the door handle or meter of the car, and play the prompt voice to remind the passengers; or the camera captures images before and after getting on the bus, and the passengers judge whether there is luggage by comparing the images Left behind.

The existing device has the problem that it cannot automatically recognize whether the passenger has opened the trunk and put the luggage in the car, cannot give targeted reminders, and the use effect is not satisfactory; in addition, the voice reminder device repeatedly plays the voice reminder, which is easy Cause fatigue and disgust of passengers and drivers.

[0005]

The search found that the Chinese utility model patent with the announcement number CN201745483U discloses a special luggage reminder device for taxis. The passenger presses the button of the device as needed when getting on the bus, and the meter starts and stops the switch or the door switch when getting off the bus. Trigger the prompt unit to send out a beep prompt and an indicator flashing prompt.

However, the patent still does not have an automatic identification function and cannot solve the above-mentioned problems.

[0006]

Utility model content

[0007]

The purpose of the utility model is to provide an automatic reminder device for luggage in the trunk of a taxi in view of the above-mentioned problems in the prior art, which can automatically determine whether passengers have opened the trunk and put their luggage in, thereby giving targeted reminders.

[0008]

In order to achieve the above objectives, the automatic reminder device for the trunk luggage of a taxi of the present invention includes a power circuit, characterized in that the output ends of the power circuit are respectively connected to the trigger circuit, the controller, and the voice circuit; the trigger circuit is connected to the backup The box-linked first switch and the second switch linked with the meter are formed, the first switch and the second switch are connected in parallel; the controller has a first monitoring input terminal connected with the first switch, and The second monitoring input terminal connected to the second switch and the output terminal connected to the voice loop.

[0009]

Further, the voice circuit includes a voice chip and a speaker respectively connected to the output terminal of the power circuit; the input terminal of the voice chip is connected to the output terminal of the controller, and the output terminal of the voice chip is connected to the speaker.

[0010]

Furthermore, the controller is a single-chip microcomputer, PLC, or industrial computer.

[0011]

Furthermore, the power circuit includes a three-terminal voltage regulator integrated circuit; the input terminal of the integrated circuit is connected to an on-board battery, and the output terminal coincides with the output terminal of the power circuit.

[0012]

The controller of the utility model judges whether the passenger has opened the trunk when boarding according to the action states of the first and second switches, and if so, it controls the voice loop to issue a voice reminder when the passenger gets off the bus.

The structure of the utility model can automatically send targeted reminders to passengers, and avoid blindly playing voice prompts repeatedly.

[0013]

Description of the drawings

[0014]

The utility model will be further explained below in conjunction with the drawings.

[0015]

Figure 1 is a schematic structural diagram of an embodiment of the utility model.

[0016]

Detailed ways

[0017]

Example

[0018]

The basic structure of the automatic reminder device for the trunk luggage of the taxi in this embodiment is shown in Figure 1, including a power circuit. The output ends of the power circuit are respectively connected to the trigger circuit, the controller 1, and the voice circuit; the trigger circuit is connected to the trunk A switch K1 is composed of a second switch K2 linked with the meter. The first switch K1 and the second switch K2 are connected in parallel; the controller 1 has a first monitoring input terminal a connected to the first switch K1 and connected to the second switch K2 The second monitoring input terminal b, and the output terminal c connected to the voice loop.

[0019]

The linkage situation of this embodiment is: when the trunk is opened, the first switch K1 is opened, when the trunk is closed, the first switch K1 is closed; when the meter starts to calculate the price, the second switch K2 is opened, and when the meter stops pricing or no pricing, the second switch is K2 is closed.

[0020]

It should be pointed out that a technical solution that is completely opposite to the linkage situation of this embodiment can also be adopted: when the trunk is opened, the first switch K1 is closed, when the trunk is closed, the first switch K1 is opened; when the meter starts to calculate the price, the second switch K2 is closed. , The second switch K2 is turned on when the meter stops pricing or does not count.

The same technical solution as in the linkage situation of this embodiment can also be adopted: the first switch K1 is opened when the trunk is opened, and the first switch K1 is closed when the trunk is closed; the second switch K2 is closed when the meter starts to count, and the meter stops pricing. Or when the price is not calculated, the second switch K2 is turned on.

Or, the first switch K1 is closed when the trunk is opened, and the first switch K1 is opened when the trunk is closed; the second switch K2 is opened when the meter starts pricing, and the second switch K2 is closed when the meter stops pricing or does not count.

These technical solutions all belong to the equivalent replacement solutions of this embodiment, and all fall within the protection scope required by the present utility model.

[0021]

The voice circuit includes a voice chip 2 (ISD1820P is used in this embodiment) and a speaker 3 respectively connected to the output of the power circuit; the input terminal d of the voice chip 2 is connected to the output terminal c of the controller, and the output terminal e is connected to the speaker 3. Connected.

[0022]

The controller 1 is a single-chip microcomputer, a PLC, or an industrial computer (the controller in this embodiment is a single-chip microcomputer, and the model is 80C51).

[0023]

The power circuit includes a three-terminal voltage stabilizing integrated circuit 4 (this embodiment uses a three-terminal voltage stabilizing integrated circuit 7805); the input terminal f of the integrated circuit 4 is connected to an on-board battery, and the output terminal g coincides with the output terminal of the power circuit.

Specifically, the power circuit of this embodiment converts the vehicle battery voltage into 5V and 3V power output, wherein the 5V power output is connected to the trigger circuit and the controller 1, and the 3V power output is connected to the voice circuit.

[0024]

The working cycle of the device in this embodiment is as follows:

[0025]

(1)Initial state: the trunk is closed, the first switch K1 is closed; the meter is not priced, the second switch K2 is closed; the first monitoring input terminal a and the second monitoring input terminal b of the controller 1 are both at high level, and the output terminal c No output; the voice loop does not work.

[0026]

(2)The passenger opens the trunk: the trunk is opened, the first switch K1 is turned on, and the first monitoring input terminal a of the controller 1 is at low level; then, the controller 1 starts to run, and the first monitoring input terminal a is in a waiting state At the same time, the state of the second monitoring input terminal b and output terminal c remains unchanged; the voice loop does not work.

[0027]

(3)The passenger closes the trunk after putting in the luggage: the trunk is closed, the first switch K1 is turned off, and the first monitoring input terminal a of the controller 1 is at a high level; then, the controller 1 sleeps the first monitoring input terminal a and enables The second monitoring input terminal b is in the waiting input state, while the state of the output terminal c remains unchanged; the voice loop does not work.

[0028]

(4)Passengers are seated and start to calculate: the meter starts to calculate, the second switch K2 is turned on, and the second monitoring input b of the controller is at low level; then, the controller 1 makes the second monitoring input b wait for input again, and at the same time The state of the first monitoring input terminal a and output terminal c remains unchanged; the voice loop does not work.

[0029]

(5)Reach the destination and stop pricing: the meter stops pricing, the second switch K2 is closed, and the second monitoring input b of the controller 1 is at a high level; then, the controller 1 sends a pulse wave to the voice loop through the output terminal c, and the voice The loop works and sends out a voice reminder to prompt passengers to take out the luggage in the trunk; at the same time, the controller 1 puts the first monitoring input terminal a in a waiting input state, and puts the second monitoring input terminal b to sleep.

[0030]

(6)The passenger opens the trunk and takes out the luggage: the trunk is opened, the first switch K1 is turned on, and the first monitoring input terminal a of the controller 1 is at a low level; then, the controller 1 makes the first monitoring input terminal a wait for input again, At the same time, the state of the second monitoring input terminal b remains unchanged, and the output terminal c has no output; the voice loop does not work.

[0031]

(7)The passenger takes out the luggage and closes the trunk: the trunk is closed, the first switch K1 is closed, and the first monitoring input a of the controller 1 is at a high level; then, the controller 1 returns to the initial state and waits for the next working cycle.

[0032]

The device of this embodiment can automatically recognize whether the passengers put their luggage in the trunk before getting on the bus, and when the passengers get off the bus, it will give targeted voice reminders to prevent the passengers from losing their items.

This embodiment has a simple structure and strong pertinence, and can avoid blindly repeating the prompt tone.

[0033]

In addition to the above implementations, the present invention can also have other implementations.

All technical solutions formed by equivalent replacements or equivalent transformations fall within the scope of protection required by the present utility model.

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

CLAIMS CN202038228U

1.

An automatic reminder for luggage in the trunk of a taxi, comprising a power circuit, characterized in that the output ends of the power circuit are respectively connected with a trigger circuit, a controller, and a voice circuit; the trigger circuit is connected by a first switch linked with the trunk It is composed of a second switch linked with a meter, the first switch and the second switch are connected in parallel; the controller has a first monitoring input terminal connected to the first switch, and is connected to the second switch The second monitoring input terminal and the output terminal connected to the voice loop.

2.

The automatic reminder device for luggage in the trunk of a taxi according to claim 1, wherein the voice circuit includes a voice chip and a speaker respectively connected to the output terminal of the power circuit; the input terminal of the voice chip and the controller The output terminal of the voice chip is connected to the speaker.

3.

The automatic reminder device for luggage in the trunk of a taxi according to claim 1 or 2, wherein the controller is a single chip microcomputer, a PLC, or an industrial computer.

4.

The automatic reminder device for luggage in the trunk of a taxi according to claim 1 or 2, wherein the power circuit includes a three-terminal voltage regulator integrated circuit; The output ends of the power supply loop coincide.

(19)中华人民共和国国家知识产权局

(12)实用新型专利

(10)授权公告号 CN 202038228 U
(45)授权公告日 2011.11.16

(21)申请号 201120113686.7

(22)申请日 2011.04.18

(73)专利权人 南京工程学院
地址 211167 江苏省南京市江宁区弘景大道1号技术转移中心

(72)发明人 刘国兵 赵又群

(74)专利代理机构 南京苏科专利代理有限责任公司 32102
代理人 何朝旭

(51)Int.Cl.
B60Q 5/00(2006.01)

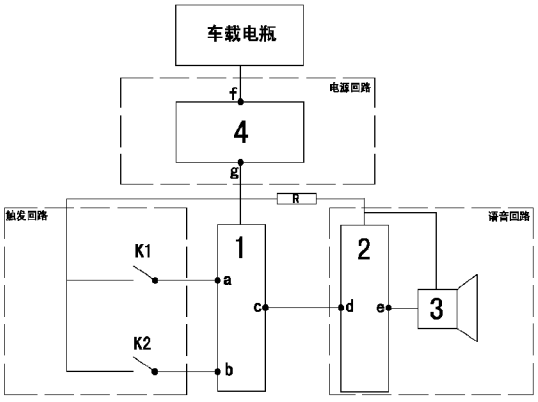
权利要求书 1 页 说明书 3 页 附图 1 页

(54)实用新型名称

出租车后备箱行李自动提醒装置

(57)摘要

本实用新型涉及一种出租车后备箱行李自动提醒装置，该装置包括电源回路，电源回路输出端分别与触发回路、控制器、语音回路相连；触发回路由与后备箱联动的第一开关和与计价器联动的第二开关构成，第一开关和第二开关并联；控制器具有与第一开关相连的第一监控输入端、与第二开关相连的第二监控输入端，以及与语音回路相连的输出端。采用本实用新型的结构可以自动向乘客发出针对性提醒，避免盲目地反复播放语音提示。



CN 202038228 U

1. 一种出租车后备箱行李自动提醒装置，包括电源回路，其特征是，所述电源回路输出端分别与触发回路、控制器、语音回路相连；所述触发回路由与后备箱联动的第一开关和与计价器联动的第二开关构成，所述第一开关和所述第二开关并联；所述控制器具有与所述第一开关相连的第一监控输入端、与所述第二开关相连的第二监控输入端，以及与语音回路相连的输出端。

2. 根据权利要求 1 所述的出租车后备箱行李自动提醒装置，其特征是，所述语音回路包括分别与所述电源回路输出端连接的语音芯片和喇叭；所述语音芯片的输入端与控制器的输出端相连，所述语音芯片的输出端与喇叭相连。

3. 根据权利要求 1 或 2 所述的出租车后备箱行李自动提醒装置，其特征是，所述控制器为单片机、PLC、或工控机。

4. 根据权利要求 1 或 2 所述的出租车后备箱行李自动提醒装置，其特征是，所述电源回路包括三端稳压集成电路；所述集成电路的输入端外接车载电瓶，其输出端与所述电源回路输出端重合。

IPR2025-00943
Tesla EX1002 Page 1397

出租车后备箱行李自动提醒装置

技术领域

[0001] 本实用新型涉及一种出租车专用的提醒装置，尤其是一种出租车后备箱行李自动提醒装置。

背景技术

[0002] 据申请人了解，乘客乘坐出租车时一般会把行李放置在汽车后备箱中，而下车时极易因忘记而造成行李遗失。现有出租车专用的提醒装置，一般通过操作汽车车门把手或计价器来开启语音电路，播放提示语音来提醒乘客；或者，通过摄像头拍摄上车前后图像，由乘客通过对比图像来判断有无行李遗落。现有装置存在的问题是，无法自动识别乘客在乘车时是否打开过后备箱放入行李，不能进行有针对性的提醒，使用效果差强人意；此外，语音提醒式的装置重复播放语音提醒，易引起乘客和驾驶员的疲劳与反感。

[0003] 检索发现，公告号 CN201745493U 的中国实用新型专利公开了一种出租车专用携带行李物品提醒装置，由乘客在上车时根据需要按下装置的按钮，在下车时由计价器启停开关或车门开关触发提示单元发出鸣叫提示和指示灯闪烁提示。但是，该专利仍然不具备自动识别功能，无法解决上述问题。

实用新型内容

[0004] 本实用新型的目的在于：针对上述现有技术存在的问题，提出一种出租车后备箱行李自动提醒装置，可自动判断乘客是否打开过后备箱放入行李，从而进行针对性提醒。

[0005] 为了达到以上目的，本实用新型的出租车后备箱行李自动提醒装置包括电源回路，其特征是，所述电源回路输出端分别与触发回路、控制器、语音回路相连；所述触发回路由与后备箱联动的第一开关和与计价器联动的第二开关构成，所述第一开关和所述第二开关并联；所述控制器具有与所述第一开关相连的第一监控输入端、与所述第二开关相连的第二监控输入端，以及与语音回路相连的输出端。

[0006] 进一步，所述语音回路包括分别与所述电源回路输出端连接的语音芯片和喇叭；所述语音芯片的输入端与控制器的输出端相连，所述语音芯片的输出端与喇叭相连。

[0007] 更进一步，所述控制器为单片机、PLC、或工控机。

[0008] 更进一步，所述电源回路包括三端稳压集成电路；所述集成电路的输入端外接车载电瓶，其输出端与所述电源回路输出端重合。

[0009] 本实用新型的控制器根据第一、第二开关的动作状态判断乘客在乘车时是否打开过后备箱，若是，则在乘客下车时控制语音回路发出语音提醒。采用本实用新型的结构可以自动向乘客发出针对性提醒，避免盲目地反复播放语音提示。

附图说明

[0010] 下面结合附图对本实用新型作进一步的说明。

[0011] 图 1 为本实用新型实施例的结构示意图。

IPR2025-00943
Tesla EX1002 Page 1398

具体实施方式

[0012] 实施例

[0013] 本实施例的出租车后备箱行李自动提醒装置基本结构如图 1 所示，包括电源回路，电源回路输出端分别与触发回路、控制器 1、语音回路相连；触发回路由与后备箱联动的第一开关 K1 和与计价器联动的第二开关 K2 构成，第一开关 K1 和第二开关 K2 并联；控制器 1 具有与第一开关 K1 相连的第一监控输入端 a、与第二开关 K2 相连的第二监控输入端 b、及与语音回路相连的输出端 c。

[0014] 本实施例的联动情况为：后备箱打开时第一开关 K1 打开，后备箱关闭时第一开关 K1 闭合；计价器开始计价时第二开关 K2 打开，计价器停止计价或未计价时第二开关 K2 闭合。

[0015] 需要指出的是，还可采用与本实施例联动情况完全相反的技术方案：后备箱打开时第一开关 K1 闭合，后备箱关闭时第一开关 K1 打开；计价器开始计价时第二开关 K2 闭合，计价器停止计价或未计价时第二开关 K2 打开。还可采用与本实施例联动情况部分相同的技术方案：后备箱打开时第一开关 K1 打开，后备箱关闭时第一开关 K1 闭合；计价器开始计价时第二开关 K2 闭合，计价器停止计价或未计价时第二开关 K2 打开。或，后备箱打开时第一开关 K1 闭合，后备箱关闭时第一开关 K1 打开；计价器开始计价时第二开关 K2 打开，计价器停止计价或未计价时第二开关 K2 闭合。这些技术方案均属于本实施例的等同替换方案，均落在本实用新型要求的保护范围。

[0016] 语音回路包括分别与所述电源回路输出端连接的语音芯片 2（本实施例采用 ISD1820P）和喇叭 3；语音芯片 2 的输入端 d 与控制器的输出端 c 相连，其输出端 e 与喇叭 3 相连。

[0017] 控制器 1 为单片机、PLC、或工控机（本实施例的控制器为单片机，型号为 80C51）。

[0018] 电源回路包括三端稳压集成电路 4（本实施例采用三端稳压集成电路 7805）；集成电路 4 的输入端 f 外接车载电瓶，其输出端 g 与电源回路输出端重合。具体而言，本实施例的电源回路将车载电瓶电压转换为 5V 和 3V 电源输出，其中 5V 电源输出分别接至触发回路和控制器 1，3V 电源输出接至语音回路。

[0019] 本实施例装置工作周期如下：

[0020] (1) 初始状态：后备箱关闭，第一开关 K1 闭合；计价器未计价，第二开关 K2 闭合；控制器 1 的第一监控输入端 a 和第二监控输入端 b 均处于高电平，输出端 c 无输出；语音回路不工作。

[0021] (2) 乘客打开后备箱：后备箱打开，第一开关 K1 打开，控制器 1 的第一监控输入端 a 处于低电平；然后，控制器 1 开始运行、并使第一监控输入端 a 处于等待输入状态，同时第二监控输入端 b、输出端 c 状态不变；语音回路不工作。

[0022] (3) 乘客放入行李后关闭后备箱：后备箱关闭，第一开关 K1 关闭，控制器 1 的第一监控输入端 a 处于高电平；然后，控制器 1 使第一监控输入端 a 休眠、并使第二监控输入端 b 处于等待输入状态，同时输出端 c 状态不变；语音回路不工作。

[0023] (4) 乘客入座、开始计价：计价器开始计价，第二开关 K2 打开，控制器的第二监控输入端 b 处于低电平；然后，控制器 1 使第二监控输入端 b 再次处于等待输入状态，同时第

IPR2025-00943
Tesla EX1002 Page 1399

一监控输入端 a、输出端 c 状态不变；语音回路不工作。

[0024] (5) 到达目的地、停止计价：计价器停止计价，第二开关 K2 闭合，控制器 1 的第二监控输入端 b 处于高电平；然后，控制器 1 经输出端 c 向语音回路发出脉冲波，语音回路工作、发出语音提醒，提示乘客取出后备箱的行李；同时，控制器 1 使第一监控输入端 a 处于等待输入状态，并使第二监控输入端 b 休眠。

[0025] (6) 乘客打开后备箱取出行李：后备箱打开，第一开关 K1 打开，控制器 1 的第一监控输入端 a 处于低电平；然后，控制器 1 使第一监控输入端 a 再次处于等待输入状态，同时第二监控输入端 b 状态不变、输出端 c 无输出；语音回路不工作。

[0026] (7) 乘客取出行李关闭后备箱：后备箱关闭，第一开关 K1 闭合，控制器 1 的第一监控输入端 a 处于高电平；然后，控制器 1 回到初始状态，等待下一工作周期。

[0027] 本实施例装置能自动识别乘客上车前，是否将行李放入后备箱中，并在乘客下车时，有针对性地发出语音提醒，防止乘客遗失物品。本实施例结构简单，针对性强，可避免盲目地重复播放提示音。

[0028] 除上述实施例外，本实用新型还可以有其他实施方式。凡采用等同替换或等效变换形成的技术方案，均落在本实用新型要求的保护范围。

IPR2025-00943
Tesla EX1002 Page 1400

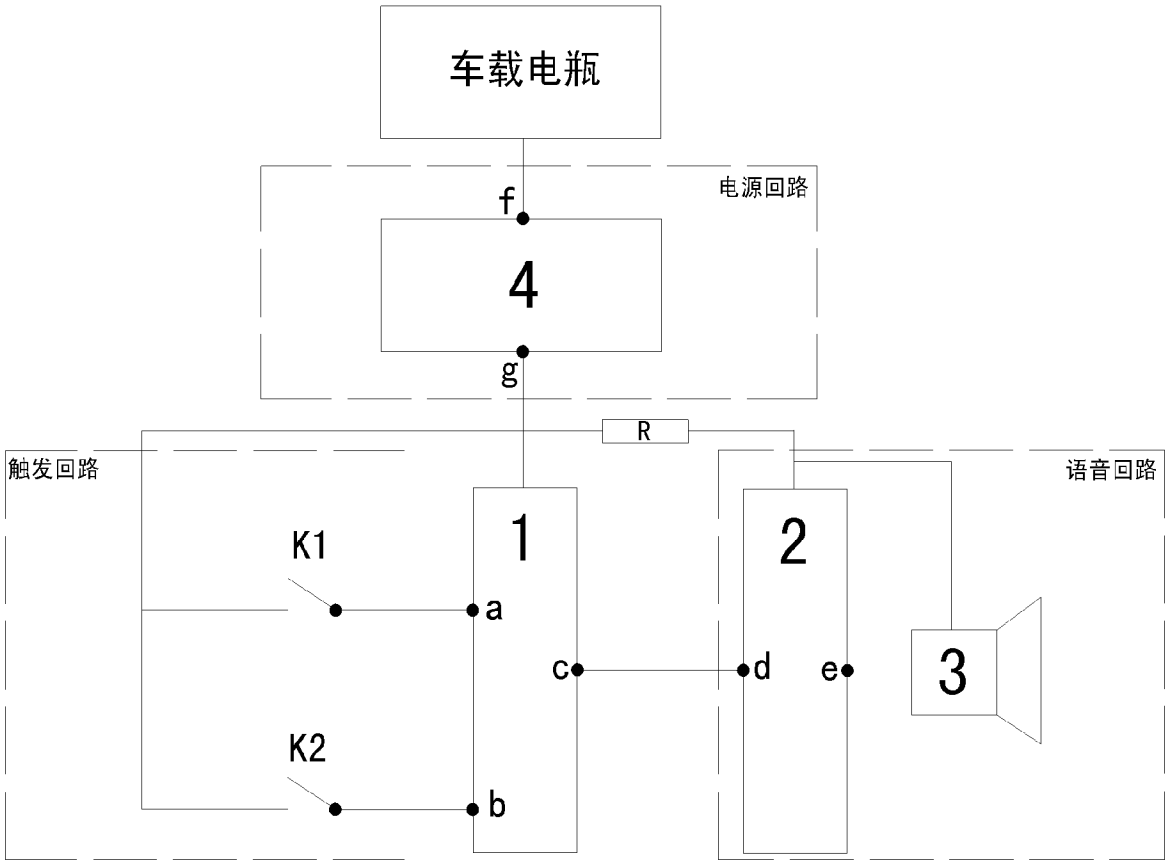


图 1

IPR2025-00943
Tesla EX1002 Page 1401

Espacenet

Bibliographic data: CN202772924 (U) — 2013-03-06

Intelligent cloud toll-gate system

Inventor(s): FANG LIANG; YE JIANYUN; XU YUBO ± (FANG LIANG, ; YE JIANYUN, ; XU YUBO)

Applicant(s): ZHEJIANG GELA WEIBAO GLASS TECHNOLOGY CO LTD ± (ZHEJIANG GELA WEIBAO GLASS TECHNOLOGY CO., LTD)

Classification: **- international:** ***H04L29/08; H04N7/18***
- cooperative:

Application number: CN201220300096U 20120621

Priority number (s): CN201220300096U 20120621

Abstract of CN202772924 (U)

Disclosed is an intelligent cloud toll-gate system, which is composed of cloud toll-gate system equipment, transmission network equipment and cloud center equipment. The cloud toll-gate system equipment includes a cloud toll-gate camera provided with a light-compensating lamp and power supply and lightning protection facilities; the transmission network equipment is used for realizing the data communication between the cloud toll-gate system equipment and the cloud center equipment, namely realizing the transmission of toll-gate data, pictures and videos; the cloud center equipment is composed of six parts including a cloud center database server, a cloud center data acquisition server, a cloud center distributed file storage server, a cloud center management application server, a cloud center communication management server and a cloud center alarm relay server. With the intelligent cloud toll-gate system, the problems of a traditional toll-gate system, such as huge amount of data, poor concurrency capability, difficultly-extended scale, long response time delay and high construction and maintenance of cost, are solved, and a technical support of high reliability, high performance and low cost is provided for large-scale and high-density deployment and multi-service and multi-maintenance application of the toll-gate system.

IPR2025-00943
Tesla EX1002 Page 1402

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

DESCRIPTION CN202772924U

Smart Cloud Bayonet System

[0001]

Technical field

[0002]

The utility model belongs to the technical field of traffic information systems, and specifically relates to an intelligent cloud bayonet system.

[0003]

Background technique

[0004]

Because the traditional electronic bayonet system completely relies on network real-time communication and central database and server for data storage, query and comparison calculations, in large-scale deployment, the computing power of network communication, central database storage and server will affect the efficiency of system operation and The stability is the bottleneck, the system scale is greatly restricted, and the construction cost is also high.

[0005]

The cloud bayonet system is the product of the close integration of cloud computing technology and bayonet system, which comprehensively utilizes distributed file system technology, distributed computing technology, high-definition image monitoring technology, Internet of Things perception detection technology, video pattern recognition technology, and wireless broadband A variety of advanced technologies such as network technology have changed the construction and application mode of traditional bayonet systems, and effectively solved the traditional bayonet system's huge data volume, poor concurrency, difficulty in scale expansion, large response delays, and high construction and maintenance costs. The problem of the formation of a cloud bayonet system with broader coverage, a larger-scale monitoring network, a larger amount of concurrent data, more business application support capabilities, and lower construction and operation costs.

It provides high-reliability, high-efficiency and low-cost technical support for the large-scale, high-density deployment and multi-service and multi-maintenance applications of the bayonet system.

[0006]

Summary of the invention

[0007]

The utility model mainly solves the technical problems existing in the above-mentioned prior art, and provides an intelligent cloud bayonet system.

[0008]

The above-mentioned technical problems of the present invention are mainly solved through the following technical solutions: a smart cloud bayonet system, which is composed of cloud bayonet system equipment, transmission network equipment and cloud center equipment, wherein the cloud bayonet system equipment is It is composed of cloud bayonet cameras with supplementary light, power supply and lightning protection facilities. The core is cloud bayonet cameras; the transmission network equipment is a high-bandwidth video cloud network, or a 3G or 4G communication network is used to form a wireless VPN mobile communication private network, or It uses WiFi wireless bridge to wirelessly bridge the cloud bayonet system equipment with the nearby optical fiber access point to realize the data communication between the cloud bayonet system equipment and the cloud center equipment, that is, to achieve bayonet data, pictures and videos The cloud center equipment is composed of the cloud center database server, the cloud center data collection server, the cloud center distributed file storage server, the cloud center management application server, the cloud center communication management server and the cloud center alarm forwarding server.

[0009]

Preferably, the cloud bayonet camera includes a high-definition CCD sensor device, an A/D conversion circuit, an audio and video encoding processing circuit, a video intelligent detection circuit, a video intelligent analysis circuit, a vehicle license plate recognition algorithm circuit, an alarm circuit, and a data storage circuit , Network communication circuit, including audio and video encoding processing circuit, video intelligent detection circuit, video intelligent analysis circuit, vehicle license plate recognition algorithm circuit, alarm circuit, data storage circuit, network communication circuit are all connected to the DSP controller, video intelligent detection circuit Together with the video intelligent analysis circuit, it controls the trigger capture circuit, the data storage circuit is connected with an expandable memory, and the DSP controller also has an external alarm control interface.

[0010]

Preferably, the cloud center database server is a server that centrally stores the passing data collected from the cloud bayonet system equipment. It adopts the Linux operating system and virtual machine architecture, and the Oracal10g database is used as the total backup storage of all bayonet data. To store cloud bayonet system equipment information, platform user information, captured vehicle information uploaded by cloud bayonet system equipment, and control alarm information, blacklisted vehicle information, and whitelisted vehicle information.

[0011]

Preferably, the cloud center data collection server adopts concurrent COP-oriented programming technology and a distributed deployment architecture to realize multi-process concurrency and automatic process monitoring, and is responsible for real-time and regular collection and reception of vehicle data captured from cloud bayonet system equipment. And store the data to the cloud bayonet system equipment and the cloud center distributed file storage server.

[0012]

Preferably, the cloud center distributed file storage server uses IP-NAS network area storage devices, supports hot-swappable disks, supports X-RAD redundancy, supports FTP file upload services, and is used to centrally save cloud bayonet system device uploads Snapshot pictures of, as evidence of motor vehicle traffic records, and also as manual authentication verification for automatic number plate recognition.

[0013]

Preferably, the cloud center management application server is responsible for the unified management of other servers, various databases, cloud bayonet system equipment and mobile terminal equipment, and provides vehicle trajectory query, deployment control, and information research and judgment application services.

[0014]

Preferably, the cloud center communication management server is responsible for real-time communication management between cloud bayonet system equipment and mobile terminal equipment, including wired communication, WIFI wireless communication, mobile phone SMS and 3G/4G wireless communication management, for data collection, vehicle control, The transmission of alarm forwarding data provides reliable communication guarantee.

[0015]

Preferably, the cloud center alarm forwarding server is responsible for forwarding the card point deployment control alarm information to the clients that need to receive alarms and key mobile phone devices. The B/S and C/S clients of the cloud center alarm forwarding server are the users and the system. The information interaction interface is the ultimate display and operation of the system business application functions.

[0016]

The utility model organically combines the Internet of Things technology, video intelligent analysis and recognition technology and cloud computing technology to form a wider coverage, a larger-scale monitoring network, a larger amount of concurrent data, more business application support capabilities and lower construction The cloud bayonet system for operation and maintenance costs is a distributed file system in cloud computing technology, and the specific application of distributed computing in vehicle bayonet systems. It effectively solves the huge data volume, poor concurrency and scale of traditional bayonet systems. Difficulty in expansion, large response time delay, and high construction and maintenance costs provide high-reliability, high-efficiency, and low-cost technical support for the large-scale, high-density deployment and multi-service and multi-maintenance applications of the bayonet system.

Compared with the traditional electronic bayonet system, this utility model has the following outstanding features: 1. It has efficient data interactive computing capabilities. Due to the use of distributed computing technology, it brings together the computing resources of all devices in the system-including Cloud bayonet system equipment and cloud center equipment, compared with the centralized computing mode of traditional electronic bayonet systems, greatly improve computing power, and provide business applications such as data query, data collision, comparison alarm, vehicle inspection, and active early warning research and judgment. Strong support; 2. Support large-scale system networking. Because the system adopts the three technologies of distributed computing, distributed file system and distributed storage in cloud computing technology, it greatly weakens the technical position of the center in the system , The operation of the system does not depend on central storage and central management, so the supporting capacity of the system is greatly improved. Conventional electronic bayonet systems generally can only support hundreds of lanes, while the scale supporting capacity of cloud bayonet systems is almost unlimited 3. With high reliability and high security, the application of cloud computing technology effectively avoids systemic collapse caused by central equipment and system failures, even if the central database server completely collapses, it will not affect vehicles within one month Business applications such as trajectory query and vehicle deployment control have greatly improved the fault tolerance to various faults, and ensure the high reliability and high safety of system operation; 4. It can quickly query trajectories and pre-warning and judging public security situation, in large-scale bayonet in the system, the cloud bayonet can provide several times the data query speed of the conventional bayonet, and collect all the front-end data information of the cloud bayonet in a short time to realize rapid warning and judgment of the public security situation; 5. Large-scale deployment and control are possible Comparison and real-time alarm. Traditional bayonet relies on database comparison when a large number of blacklisted vehicles and a wide range of bayonet control are deployed. The response time is slow and the timeliness is poor. The cloud bayonet system directly sends the deployment control instructions and information to the front end. The equipment is compared by each lane by itself, which greatly improves the comparison efficiency and ensures the real-time performance of the deployment and control alarm; 6. The construction cost is low, and it is suitable for various on-site environments. The cloud bayonet system can not rely on optical fiber networking. , The captured data can be returned in real time without real-time return, and it can be used to interact with the application on the 3G/4G network, and it also supports full video capture triggering. There is no need to place a coil on the road, which greatly saves the cost of network leasing and on-site construction. At the same time, it supports solar power supply and greatly reduces the dependence on the on-site environment, making it more suitable for edge and complex monitoring sites.

[0017]

Description of the drawings

[0018]

Figure 1 is a schematic diagram of the system structure of the utility model;

[0019]

Figure 2 is a schematic diagram of a principle structure of the cloud bayonet camera of the present invention.

[0020]

Detailed ways

[0021]

In the following, the technical solution of the present utility model will be further described in detail through the embodiments and in conjunction with the drawings.

[0022]

Embodiment: Referring to Figure 1, Figure 2 and Figure 3, this utility model is composed of cloud bayonet system equipment, transmission network equipment and cloud center equipment, among which:

[0023]

The cloud bayonet system equipment consists of cloud bayonet cameras with supplementary light, power supply and lightning protection facilities. The core is the cloud bayonet camera. The cloud bayonet camera includes high-definition CCD sensor components, A/D conversion circuit, audio and video encoding processing Circuit, video intelligent detection circuit, video intelligent analysis circuit, vehicle number plate recognition algorithm circuit, alarm circuit, data storage circuit, network communication circuit, including audio and video encoding processing circuit, video intelligent detection circuit, video intelligent analysis circuit, vehicle number plate The recognition algorithm circuit, alarm circuit, data storage circuit, and network communication circuit are all connected to the DSP controller. The video intelligent detection circuit and the video intelligent analysis circuit jointly control the trigger capture circuit. The data storage circuit is connected with an expandable memory. The DSP controller also Equipped with external alarm control interface;

[0024]

The transmission network equipment is a high-bandwidth video cloud network, or a 3G or 4G communication network is used to form a wireless VPN mobile communication private network, or a WiFi wireless bridge is used to wirelessly bridge the cloud bayonet system equipment with the nearby optical fiber access point. It is used to realize the data communication between the cloud bayonet system equipment and the cloud center equipment, that is, to realize the transmission of bayonet data, pictures and videos;

[0025]

The cloud center equipment is composed of the cloud center database server, the cloud center data collection server, the cloud center distributed file storage server, the cloud center management application server, the cloud center communication management server, and the cloud center alarm forwarding server:

[0026]

The cloud center database server is a server that centrally stores the passing data collected from the cloud bayonet system equipment. It uses the Linux operating system and virtual machine architecture, and the Oracal10g database is used as the total backup storage for all bayonet data to store the cloud bayonet System equipment information, platform user information, captured vehicle information uploaded by cloud bayonet system equipment, deployment control alarm information, blacklisted vehicle information, and whitelisted vehicle information;

[0027]

The cloud center data collection server adopts concurrent COP-oriented programming technology and a distributed deployment architecture to realize multi-process concurrency and automatic process monitoring. It is responsible for real-time and regular collection and reception of vehicle data captured from cloud bayonet system equipment, and stores the data in Cloud bayonet system equipment and cloud center distributed file storage server;

[0028]

The cloud center distributed file storage server uses IP-NAS network area storage devices, supports hot-swappable disks, supports X-RAD redundancy, supports FTP file upload services, and is used to centrally save snapshots uploaded by cloud bayonet system devices. The evidence of motor vehicle traffic records is also used as a manual identification verification for automatic number plate recognition;

[0029]

The cloud center management application server is responsible for the unified management of other servers, databases, cloud bayonet system equipment and mobile terminal equipment, and provides vehicle trajectory query, deployment control, and information research and judgment application services;

[0030]

The cloud center communication management server is responsible for real-time communication management between cloud bayonet system equipment and mobile terminal equipment, including wired communication, WIFI wireless communication, mobile phone SMS and 3G/4G wireless communication management, and data transmission for data collection, vehicle control, and alarm forwarding Provide reliable communication guarantee;

[0031]

The cloud center alarm forwarding server is responsible for forwarding the card point deployment control alarm information to the clients that need to receive alarms and key mobile devices. The B/S and C/S clients of the cloud center alarm forwarding server are the information interaction interface between the user and the system. It is the final display and operation of the system's business application functions.

[0032]

The utility model organically combines the Internet of Things technology, video intelligent analysis and recognition technology and cloud computing technology to form a wider coverage, a larger-scale monitoring network, a larger amount of concurrent data, more business application support capabilities and lower construction The cloud bayonet system for operation and maintenance costs is a distributed file system in cloud computing technology, and the specific application of distributed computing in vehicle bayonet systems. It effectively solves the huge data volume, poor concurrency and scale of traditional bayonet systems. Difficulty in expansion, large response time delay, and high construction and maintenance costs provide high-reliability, high-efficiency, and low-cost technical support for the large-scale, high-density deployment and multi-service and multi-maintenance applications of the bayonet system.

[0033]

Finally, it should be pointed out that the above embodiments are only representative examples of the present invention.

Obviously, the present invention is not limited to the above-mentioned embodiments, and many variations are possible.

IPR2025-00943
Tesla EX1002 Page 1407

Any simple modifications, equivalent changes and modifications made to the above embodiments based on the technical essence of the present utility model shall be deemed to belong to the protection scope of the present utility model.

IPR2025-00943
Tesla EX1002 Page 1408

Patent Translate
Powered by EPO and Google

Notice

This translation is machine-generated. It cannot be guaranteed that it is intelligible, accurate, complete, reliable or fit for specific purposes. Critical decisions, such as commercially relevant or financial decisions, should not be based on machine-translation output.

CLAIMS CN202772924U

1.

A smart cloud bayonet system, characterized in that the system is composed of cloud bayonet system equipment (1), transmission network equipment (2) and cloud center equipment (3), wherein the cloud bayonet system equipment (1) consists of It is composed of cloud bayonet cameras (11) for fill light, power supply and lightning protection facilities, the core is cloud bayonet cameras; the transmission network equipment (2) is a high-bandwidth video cloud network, or a 3G or 4G communication network is used to form a wireless VPN Mobile communication private network, or use WiFi wireless bridge to wirelessly bridge the cloud bayonet system equipment with the nearby optical fiber access point to realize the data communication between the cloud bayonet system equipment and the cloud center equipment, that is, to realize the card Transmission of port data, pictures and videos; cloud center equipment (3) consists of cloud center database server (31), cloud center data acquisition server (32), cloud center distributed file storage server (33), cloud center management application server (34), the cloud center communication management server (35) and the cloud center alarm forwarding server (36) are composed of six parts.

2.

The smart cloud bayonet system according to claim 1, wherein the cloud bayonet camera includes a high-definition CCD sensor device, an A/D conversion circuit, an audio and video encoding processing circuit, a video intelligent detection circuit, a video intelligent analysis circuit, Vehicle number plate recognition algorithm circuit, alarm circuit, data storage circuit, network communication circuit, including audio and video encoding processing circuit, video intelligent detection circuit, video intelligent analysis circuit, vehicle number plate recognition algorithm circuit, alarm circuit, data storage circuit, network The communication circuits are all connected to the DSP controller, the video intelligent detection circuit and the video intelligent analysis circuit jointly control the triggering capture circuit, the data storage circuit is connected with an expandable memory, and the DSP controller also has an external alarm control interface.

3.

The smart cloud bayonet system according to claim 1, characterized in that the cloud center database server (31) is a server that centrally stores the passing data collected from the cloud bayonet system equipment, using Linux operating system and virtual Machine architecture, Oracal10g database is used as the total backup storage of all bayonet data, used to store cloud bayonet system device information, platform user information, cloud bayonet system equipment uploading captured vehicle information and deployment control alarm information, blacklisted vehicle information, Whitelist vehicle information.

4.

The smart cloud bayonet system according to claim 1, characterized in that the cloud center data collection server (32) adopts a concurrent COP-oriented programming technology, a distributed deployment architecture, realizes multi-process concurrency

and automatic process monitoring, and is responsible for Cloud bayonet system equipment (1) capture and receive real-time and timing data of passing cars, and store the data in the cloud bayonet system equipment (1) and cloud center distributed file storage server (33).

5.

The smart cloud bayonet system according to claim 1, wherein the cloud center distributed file storage server (33) selects an IP-NAS network area storage device, supports hot-swappable disks, and supports X-RAD redundancy, Support FTP file upload service, used to centrally save the snapshot pictures uploaded by the cloud bayonet system equipment (1), as evidence of motor vehicle traffic records, and also as a manual authentication verification for automatic number plate recognition.

6.

The smart cloud bayonet system according to claim 1, characterized in that the cloud center management application server (34) is responsible for the unified management of other servers, databases, cloud bayonet system equipment and mobile terminal equipment, and provides vehicles Track query, deployment control, information research and judgment application services.

7.

The smart cloud bayonet system according to claim 1, characterized in that the cloud center communication management server (35) is responsible for real-time communication management between cloud bayonet system equipment and mobile terminal equipment, including wired communication, WIFI wireless communication, and mobile phone SMS and 3G/4G wireless communication management provide reliable communication guarantee for data collection, vehicle deployment control, and alarm forwarding data transmission.

8.

The smart cloud bayonet system according to claim 1, characterized in that the cloud center alarm forwarding server (36) is responsible for forwarding the alarm information of the card point deployment control to the client and the key device of the mobile phone that need to receive the alarm, and the B/ S, C/S client is the information interaction interface between the user and the system, and is the final display and operation of the system's business application functions.

IPR2025-00943
Tesla EX1002 Page 1410

(19)中华人民共和国国家知识产权局

(12)实用新型专利

(10)授权公告号 CN 202772924 U
(45)授权公告日 2013.03.06

(21)申请号 201220300096.X

(22)申请日 2012.06.21

(73)专利权人 浙江广信智能建筑研究院有限公司
地址 310004 浙江省杭州市费家塘路576号6楼

(72)发明人 方良 叶建云 徐玉波

(74)专利代理机构 浙江永鼎律师事务所 33233
代理人 陆永强

(51)Int.Cl.
H04L 29/08(2006.01)
H04N 7/18(2006.01)

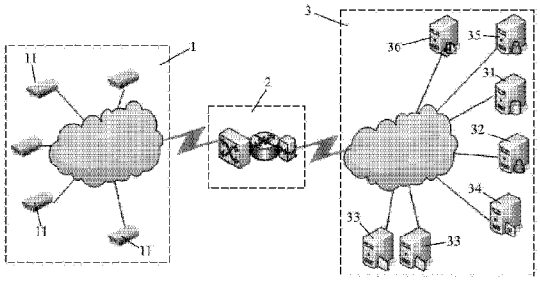
权利要求书 1 页 说明书 4 页 附图 1 页

(54)实用新型名称

智能云卡口系统

(57)摘要

一种智能云卡口系统，由云端卡口系统设备、传输网络设备和云中心设备组成，其中云端卡口系统设备由带补光灯、供电和防雷设施的云端卡口摄像机组成；传输网络设备实现云端卡口系统设备与云中心设备之间的数据通信，即实现卡口数据、图片和视频的传输；云中心设备由云中心数据库服务器、云中心数据采集服务器、云中心分布式文件存储服务器、云中心管理应用服务器、云中心通信管理服务器和云中心报警转发服务器六部分组成。本实用新型有效的解决了传统卡口系统数据量庞大、并发能力差、规模难扩展、响应时延大、建设和维护成本高的难题，为卡口系统的大规模、高密度部署和多业务、多维护应用提供了高可靠、高效能、低成本的技术支撑。



CN 202772924 U

IPR2025-00943
Tesla EX1002 Page 1411

1. 一种智能云卡口系统，其特征在于所述系统由云端卡口系统设备(1)、传输网络设备(2)和云中心设备(3)组成，其中云端卡口系统设备(1)由带补光灯、供电和防雷设施的云端卡口摄像机(11)组成，核心是云端卡口摄像机；传输网络设备(2)是高宽带视频云网络，或者是采用 3G 或 4G 通信网络组成无线 VPN 移动通信专网，或者是利用 WiFi 无线网桥将云端卡口系统设备与就近的光纤接入点进行无线桥接，用以实现云端卡口系统设备与云中心设备之间的数据通信，即实现卡口数据、图片和视频的传输；云中心设备(3)由云中心数据库服务器(31)、云中心数据采集服务器(32)、云中心分布式文件存储服务器(33)、云中心管理应用服务器(34)、云中心通信管理服务器(35)和云中心报警转发服务器(36)六部分组成。

2. 根据权利要求 1 所述的智能云卡口系统，其特征在于所述云端卡口摄像机包括高清 CCD 传感器件、A/D 转换电路、音视频编码处理电路、视频智能检测电路、视频智能分析电路、车辆号牌识别算法电路、报警电路、数据存储电路、网络通信电路，其中音视频编码处理电路、视频智能检测电路、视频智能分析电路、车辆号牌识别算法电路、报警电路、数据存储电路、网络通信电路均连接至 DSP 控制器，视频智能检测电路和视频智能分析电路共同控制触发抓拍电路，数据存储电路上连接有可扩展存储器，DSP 控制器还设有外部报警控制接口。

3. 根据权利要求 1 所述的智能云卡口系统，其特征在于所述云中心数据库服务器(31)是对从云端卡口系统设备采集的过车数据进行集中存储的服务器，采用 Linux 操作系统和虚拟机架构、Oracal10g 数据库作为所有卡口数据的总备份存储，用来存储云端卡口系统设备信息、平台用户信息、云端卡口系统设备上传的抓拍过车信息和布控报警信息、黑名单车辆信息、白名单车辆信息。

4. 根据权利要求 1 所述的智能云卡口系统，其特征在于所述云中心数据采集服务器(32)采用面向并发 COP 的编程技术，分布式部署架构，实现多进程并发和进程自动监控，负责从云端卡口系统设备(1)抓拍过车数据的实时、定时采集接收，并将数据存储到云端卡口系统设备(1)和云中心分布式文件存储服务器(33)。

5. 根据权利要求 1 所述的智能云卡口系统，其特征在于所述云中心分布式文件存储服务器(33)选用 IP-NAS 网络区域存储设备，支持磁盘热插拔，支持 X-RAD 冗余，支持 FTP 文件上传服务，用于集中保存云端卡口系统设备(1)上传的抓拍图片，作为机动车通行记录的佐证，同时也作为对自动号牌识别的人工鉴别验证。

6. 根据权利要求 1 所述的智能云卡口系统，其特征在于所述云中心管理应用服务器(34)负责对其他服务器、各数据库、云端卡口系统设备和移动终端设备的统一管理，并提供车辆轨迹查询、布控、信息研判应用服务。

7. 根据权利要求 1 所述的智能云卡口系统，其特征在于所述云中心通信管理服务器(35)负责云端卡口系统设备与移动终端设备的实时通信管理，包括有线通信、WIFI 无线通信、手机短信和 3G/4G 无线通信管理，为数据采集、车辆布控、报警转发数据的传输提供可靠的通信保障。

8. 根据权利要求 1 所述的智能云卡口系统，其特征在于所述云中心报警转发服务器(36)负责对卡点布控报警信息转发给需要接收报警的客户端、手机重点设备上，其 B/S、C/S 客户端是用户与系统的信息交互界面，是系统业务应用功能的最终展现和操作。

IPR2025-00943
Tesla EX1002 Page 1412

智能云卡口系统

技术领域

[0001] 本实用新型属于交通信息系统技术领域，具体是涉及一种智能云卡口系统。

背景技术

[0002] 由于传统的电子卡口系统完全依赖网络实时通信与中心数据库和服务器进行数据存储、查询和比对运算，在大规模部署时，网络通信、中心数据库存储和服务器的运算能力成为影响系统运行效率和稳定性的瓶颈，系统规模受到很大的制约，建设成本也高。

[0003] 云卡口系统是云计算技术与卡口系统紧密结合的产物，综合利用了分布式文件系统技术、分布式计算技术、高清图像监控技术、物联网感知探测技术、视频模式识别技术、有无线宽带网络技术等多种先进技术，改变了传统卡口系统的建设和应用模式，有效的解决了传统卡口系统数据量庞大、并发能力差、规模难扩展、响应时延大、建设和维护成本高的难题，形成了具有更广阔覆盖面、更大规模监控网络、更大并发数据量、更多业务应用支撑能力和更低建设运维成本的云卡口系统。为卡口系统的大规模、高密度部署和多业务、多维护应用提供了高可靠、高效能、低成本的技术支撑。

发明内容

[0004] 本实用新型主要是解决上述现有技术所存在的技术问题，提供了一种智能云卡口系统。

[0005] 本实用新型的上述技术问题主要是通过下述技术方案得以解决的：一种智能云卡口系统，由云端卡口系统设备、传输网络设备和云中心设备组成，其中云端卡口系统设备由带补光灯、供电和防雷设施的云端卡口摄像机组成，核心是云端卡口摄像机；传输网络设备是高宽带视频云网络，或者是采用 3G 或 4G 通信网络组成无线 VPN 移动通信专网，或者是利用 WiFi 无线网桥将云端卡口系统设备与就近的光纤接入点进行无线桥接，用以实现云端卡口系统设备与云中心设备之间的数据通信，即实现卡口数据、图片和视频的传输；云中心设备由云中心数据库服务器、云中心数据采集服务器、云中心分布式文件存储服务器、云中心管理应用服务器、云中心通信管理服务器和云中心报警转发服务器六部分组成。

[0006] 作为优选，所述云端卡口摄像机包括高清 CCD 传感器件、A/D 转换电路、音视频编码处理电路、视频智能检测电路、视频智能分析电路、车辆号牌识别算法电路、报警电路、数据存储电路、网络通信电路，其中音视频编码处理电路、视频智能检测电路、视频智能分析电路、车辆号牌识别算法电路、报警电路、数据存储电路、网络通信电路均连接至 DSP 控制器，视频智能检测电路和视频智能分析电路共同控制触发抓拍电路，数据存储电路上连接有可扩展存储器，DSP 控制器还设有外部报警控制接口。

[0007] 作为优选，所述云中心数据库服务器是对从云端卡口系统设备采集的过车数据进行集中存储的服务器，采用 Linux 操作系统和虚拟机架构、Oracal10g 数据库作为所有卡口数据的总备份存储，用来存储云端卡口系统设备信息、平台用户信息、云端卡口系统设备上传的抓拍过车信息和布控报警信息、黑名单车辆信息、白名单车辆信息。

IPR2025-00943
Tesla EX1002 Page 1413

[0008] 作为优选，所述云中心数据采集服务器采用面向并发 COP 的编程技术，分布式部署架构，实现多进程并发和进程自动监控，负责从云端卡口系统设备抓拍过车数据的实时、定时采集接收，并将数据存储到云端卡口系统设备和云中心分布式文件存储服务器。

[0009] 作为优选，所述云中心分布式文件存储服务器选用 IP-NAS 网络区域存储设备，支持磁盘热插拔，支持 X-RAD 冗余，支持 FTP 文件上传服务，用于集中保存云端卡口系统设备上传的抓拍图片，作为机动车通行记录的佐证，同时也作为对自动号牌识别的人工鉴别验证。

[0010] 作为优选，所述云中心管理应用服务器负责对其他服务器、各数据库、云端卡口系统设备和移动终端设备的统一管理，并提供车辆轨迹查询、布控、信息研判应用服务。

[0011] 作为优选，所述云中心通信管理服务器负责云端卡口系统设备与移动终端设备的实时通信管理，包括有线通信、WIFI 无线通信、手机短信和 3G/4G 无线通信管理，为数据采集、车辆布控、报警转发数据的传输提供可靠的通信保障。

[0012] 作为优选，所述云中心报警转发服务器负责对卡点布控报警信息转发给需要接收报警的客户端、手机重点设备上，云中心报警转发服务器的 B/S、C/S 客户端是用户与系统的信息交互界面，是系统业务应用功能的最终展现和操作。

[0013] 本实用新型将物联网技术、视频智能分析识别技术和云计算技术的有机结合，形成了具有更广阔覆盖面、更大规模监控网络、更大并发数据量、更多业务应用支撑能力和更低建设运维成本的云卡口系统，是云计算技术中的分布式文件系统、分布式计算在车辆卡口系统中的具体应用，有效的解决了传统卡口系统数据量庞大、并发能力差、规模难扩展、响应时延大、建设和维护成本高的难题，为卡口系统的大规模、高密度部署和多业务、多维护应用提供了高可靠、高效能、低成本的技术支撑。相对于传统电子卡口系统，本实用新型具有以下几个方面的突出特点：1、具有高效的数据交互运算能力，由于采用了分布式计算技术，汇集了系统内所有设备的计算资源——包括云端卡口系统设备和云中心设备，与传统电子卡口系统的集中运算模式相比，大大提高了运算能力，为数据查询、数据碰撞、比对报警、车辆排查、主动预警研判等业务应用提供强有力的支撑；2、支持大规模系统联网，由于系统采用了云计算技术中的分布式计算、分布式文件系统和分布式存储三大技术，极大的弱化了中心在系统中的技术地位，系统的运行不依赖于中心存储和中心管理，因此对系统规模的支撑能力大大提高，常规电子卡口系统一般只能支持数百个车道规模，而云卡口系统的规模支持能力几乎是无限的；3、具有高可靠和高安全性，云计算技术的应用有效的避免了因中心设备、系统故障而导致的系统性崩溃，即便是中心数据库服务器完全崩溃，也不影响一个月以内的车辆轨迹查询和车辆布控等业务应用，大大提高了对各种故障的容错能力，保证了系统运行的高可靠性和高安全性；4、可以轨迹快速查询和治安态势预警研判，在大规模卡口系统中，云卡口能够提供数倍于常规卡口的数据查询速度，在很短的时间内汇集所有云卡口前端的数据信息，实现治安态势的快速预警和研判；5、可以大规模布控比对和实时报警，传统卡口在大量黑名单车辆、大范围卡口布控时，依赖于数据库的比对，反应时间慢，时效性差，云卡口系统则直接将布控指令和信息发送到前端设备，由每个车道自行比对，极大的提高了比对效率，保证了布控报警的实时性；6、建设成本低，适应于各种现场环境，云卡口系统可以不依赖于光纤联网，抓拍数据可以不用实时回传，完全可以借助于 3G/4G 网络与应用端进行业务交互，而且还支持完全视频捕捉触发，无需在道路上安置线

IPR2025-00943
Tesla EX1002 Page 1414

圈，人人节省了网络租赁和现场施工的费用，同时支持太阳能供电，对现场环境的依赖也大大降低，更适合于边缘、复杂的监控现场。

附图说明

[0014] 图 1 是本实用新型的一种系统结构示意图；

[0015] 图 2 是本实用新型云端卡口摄像机的一种原理结构示意图。

具体实施方式

[0016] 下面通过实施例，并结合附图，对本实用新型的技术方案作进一步具体的说明。

[0017] 实施例：参看图 1、图 2 和图 3，本实用新型由云端卡口系统设备、传输网络设备和云中心设备组成，其中：

[0018] 云端卡口系统设备由带补光灯、供电和防雷设施的云端卡口摄像机组成，核心是云端卡口摄像机，云端卡口摄像机包括高清 CCD 传感器件、A/D 转换电路、音视频编码处理电路、视频智能检测电路、视频智能分析电路、车辆号牌识别算法电路、报警电路、数据存储电路、网络通信电路，其中音视频编码处理电路、视频智能检测电路、视频智能分析电路、车辆号牌识别算法电路、报警电路、数据存储电路、网络通信电路均连接至 DSP 控制器，视频智能检测电路和视频智能分析电路共同控制触发抓拍电路，数据存储电路上连接有可扩展存储器，DSP 控制器还设有外部报警控制接口；

[0019] 传输网络设备是高宽带视频云网络，或者是采用 3G 或 4G 通信网络组成无线 VPN 移动通信专网，或者是利用 WiFi 无线网桥将云端卡口系统设备与就近的光纤接入点进行无线桥接，用以实现云端卡口系统设备与云中心设备之间的数据通信，即实现卡口数据、图片和视频的传输；

[0020] 云中心设备由云中心数据库服务器、云中心数据采集服务器、云中心分布式文件存储服务器、云中心管理应用服务器、云中心通信管理服务器和云中心报警转发服务器六部分组成：

[0021] 云中心数据库服务器是对从云端卡口系统设备采集的过车数据进行集中存储的服务器，采用 Linux 操作系统和虚拟机架构、Oracal10g 数据库作为所有卡口数据的总备份存储，用来存储云端卡口系统设备信息、平台用户信息、云端卡口系统设备上传的抓拍过车信息和布控报警信息、黑名单车辆信息、白名单车辆信息；

[0022] 云中心数据采集服务器采用面向并发 COP 的编程技术，分布式部署架构，实现多进程并发和进程自动监控，负责从云端卡口系统设备抓拍过车数据的实时、定时采集接收，并将数据存储到云端卡口系统设备和云中心分布式文件存储服务器；

[0023] 云中心分布式文件存储服务器选用 IP-NAS 网络区域存储设备，支持磁盘热插拔，支持 X-RAD 冗余，支持 FTP 文件上传服务，用于集中保存云端卡口系统设备上传的抓拍图片，作为机动车通行记录的佐证，同时也作为对自动号牌识别的人工鉴别验证；

[0024] 云中心管理应用服务器负责对其他服务器、各数据库、云端卡口系统设备和移动终端设备的统一管理，并提供车辆轨迹查询、布控、信息研判应用服务；

[0025] 云中心通信管理服务器负责云端卡口系统设备与移动终端设备的实时通信管理，包括有线通信、WIFI 无线通信、手机短信和 3G/4G 无线通信管理，为数据采集、车辆布控、报

IPR2025-00943
Tesla EX1002 Page 1415

警转发数据的传输提供可靠的通信保障；

[0026] 云中心报警转发服务器负责对卡点布控报警信息转发给需要接收报警的客户端、手机重点设备上，云中心报警转发服务器的 B/S、C/S 客户端是用户与系统的信息交互界面，是系统业务应用功能的最终展现和操作。

[0027] 本实用新型将物联网技术、视频智能分析识别技术和云计算技术的有机结合，形成了具有更广阔覆盖面、更大规模监控网络、更大并发数据量、更多业务应用支撑能力和更低建设运维成本的云卡口系统，是云计算技术中的分布式文件系统、分布式计算在车辆卡口系统中的具体应用，有效的解决了传统卡口系统数据量庞大、并发能力差、规模难扩展、响应时延大、建设和维护成本高的难题，为卡口系统的大规模、高密度部署和多业务、多维护应用提供了高可靠、高效能、低成本的技术支撑。

[0028] 最后，应当指出，以上实施例仅是本实用新型较有代表性的例子。显然，本实用新型不限于上述实施例，还可以有许多变形。凡是依据本实用新型的技术实质对以上实施例所作的任何简单修改、等同变化与修饰，均应认为属于本实用新型的保护范围。

IPR2025-00943
Tesla EX1002 Page 1416

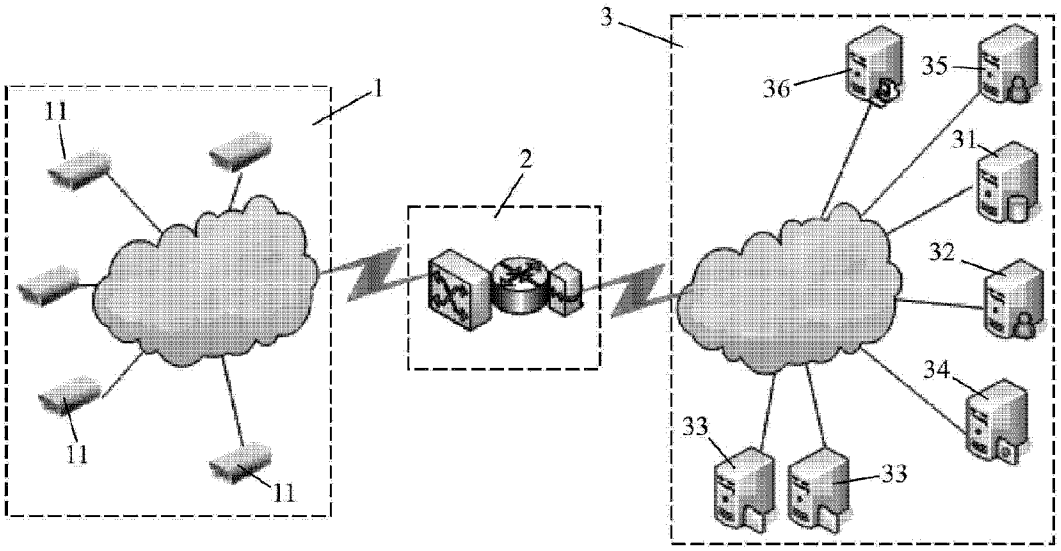


图 1

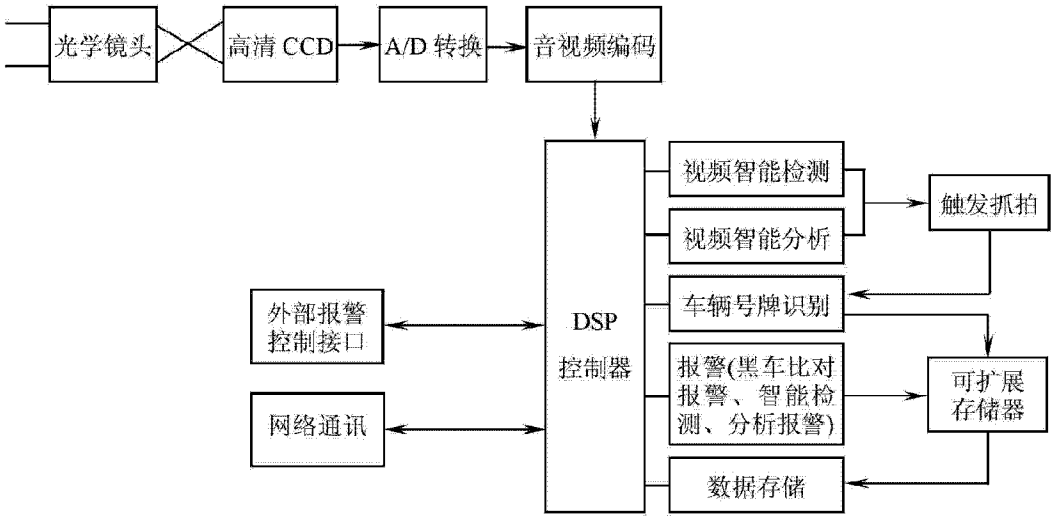


图 2

IPR2025-00943
Tesla EX1002 Page 1417

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 582 236 B1

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
03.06.1998 Bulletin 1998/23

(51) Int. Cl.[6]: **B60R 1/00**, B60R 21/00, B60K 28/00, G05D 1/02

(21) Application number: **93112302.0**

(22) Date of filing: **30.07.1993**

(54) **Vehicle crash predictive and evasive operation system by neural networks**

Steuerungssystem mit neuronalem Netzwerk zum Voraussagen und Vermeiden von Zusammenstössen eines Fahrzeuges

Système de commande à réseau neuronal pour prévoir et éviter les collisions pour un véhicule

(84) Designated Contracting States:
DE GB

(30) Priority: **04.08.1992 JP 229201/92**

(43) Date of publication of application:
09.02.1994 Bulletin 1994/06

(73) Proprietor: **TAKATA CORPORATION**
Minato-ku, Tokyo 106 (JP)

(72) Inventor: **Nishio, Tomoyuki**
Kawasaki-shi, Kanagawa-ken (JP)

(74) Representative:
Heim, Hans-Karl, Dipl.-Ing. et al
Weber & Heim
Patentanwälte
Irmgardstrasse 3
81479 München (DE)

(56) References cited:
EP-A- 0 358 628
WO-A-90/02985
DE-A- 3 837 054
DE-A- 4 001 493
FR-A- 2 554 612

- **PATENT ABSTRACTS OF JAPAN vol. 16, no. 157 (M-1236)16 April 1992**
- **RUMELHART ET AL 'Parallel Distributed Processing, vol. 1: Foundations' 1986 , THE MIT PRESS , CAMBRIDGE, MASSACHUSETTS, USA**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 0 582 236 B1

Printed by Xerox (UK) Business Services
2.16.3/3.4

IPR2025-00943
Tesla EX1002 Page 1418

Description

Background of the Invention

The invention relates to an apparatus for predicting and evading potential crash of a vehicle and for predicting real crash of a vehicle and acting correspondingly according to the preamble of claim 1. Furthermore, the invention relates to a method according to the preamble of claim 9.

A driver has an unconscious and immediate sense of various conditions through the objects in view and, as a case may be, he must take an action to evade any possible crash or collision. However, drivers will often be panicked at the emergency of above their sense. Such a panicked driver may sometimes be the last one who can cope with the emergency to ensure the active safety of the vehicle. Besides, the response delay to stimuli in varying degrees is inherent to human beings, so that it is impossible in some cases to evade crash or danger by physical considerations. With this respect, various techniques have been developed to evade collision by means of mounting on a vehicle a system for determining the possibility of crash in a mechanical or electrical manner before it happens. Accidents could be reduced if drivers had an automatic system or the like warning of potential collision situations.

An apparatus comprising a system for predicting and evading crash of a vehicle including image pick-up means, crash predicting means, safety drive ensuring means and a neural network according to the preamble of claim 1 is known from JP-A-4008639 (Abstract).

In WO-A-90/02985 a system for predicting and evading crash of a vehicle comprising a sensor system mounted on the vehicle for observing the surrounding in a running direction is disclosed. A crash predicting system is successively supplied with actual sensor system data for predicting occurence of crash between the vehicle and potential dangerous objects on the roadway and for producing an operational signal when there is possibility of crash. In response to the operational signal the vehicle is operated to evade the crash.

Another moveable device with supersonic sensors and a neural network to further process the data taken by the supersonic sensor system is disclosed in DE-A-4001493. Furthermore, in this moveable device learning data are used to train the neural network.

In general, neural networks are known from "Parallel distributed Processing", Vol. 1: Foundations, A Bradford Book, The MIT Press, Cambridge, Mass., London, England, by David E. Rumelhart, James L. McClelland and the PDP Research Group.

An automobile collision avoidance radar is typically used as this automatic system. Such an automobile collision avoidance radar is disclosed in, for example, M. Kiyoto and A. Tachibana, Nissan Technical Review: Automobile Collision-Avoidance Radar, Vol. 18, Dec. 1982 that is incorporated by reference herein in its entirety. The radar disclosed comprises a small radar radiation element and antennas installed at the front end of a vehicle. A transmitter transmits microwaves through the radiation element towards the headway. The microwave backscatter from a leading vehicle or any other objects as echo returns. The echo returns are received by a receiver through the antennas and supplied to a signal processor. The signal processor carries out signal processing operation to calculate a relative velocity and a relative distance between the object and the vehicle. The relative velocity and the relative distance are compared with predetermined values, respectively, to determine if the vehicle is going to collide with the object. The high possibility of collision results in activation of a proper safety system or systems.

However, the above mentioned radar system has a disadvantage of faulty operation or malfunctions, especially when the vehicle implementing this system passes by a sharp curve in a road. The radar essentially detects objects in front of the vehicle on which it is mounted. The system thus tends to incorrectly identify objects alongside the road such as a roadside, guard rails or even an automobile correctly running on the adjacent lane.

An intelligent vehicle has also been proposed that comprises an image processing system for cruise and traction controls. Ever-changing views spreading ahead the vehicle are successively picked up as image patterns. These image patterns are subjected to pattern matching with predetermined reference patterns. The reference patterns are classified into some categories associated with possible driving conditions. For example, three categories are defined for straight running, right turn and left turn. When a matching result indicates the presence of potentially dangerous objects in the picked up image, a steering wheel and a brake system are automatically operated through a particular mechanism to avoid or evade crash to that object.

The image processing system of the type described is useful in normal driving conditions where the pattern matching can be effectively made between the image patterns successively picked up and the reference patterns for safety driving control. However, image patterns representing various conditions on the roadway should be stored previously in the intelligent vehicle as the reference patterns. Vehicle orientation at initiation of crash varies greatly, so that huge numbers of reference patterns are required for the positive operation. This means that only a time-consuming calculation will result in a correct matching of the patterns, which is not suitable for evading an unexpected crash.

It is, of course, possible to increase operational speed of the pattern matching by using a large dedicated image processor. However, such a dedicated processor is generally complex in structure and relatively expensive, so that it is difficult to apply the same as the on-vehicle equipment. In addition, on-vehicle image processors, if achieved, will perform its function

IPR2025-00943
Tesla EX1002 Page 1419

sufficiently only in the limited applications such as a supplemental navigation system during the normal cruising.

Summary of the Invention

An object of the present invention is to provide a system for predicting and evading crash of a vehicle using neural networks.

Another object of the present invention is to provide a system capable of training neural networks by means of collecting image data representing ever-changing vistas along the travel direction of a vehicle until the vehicle collides with something.

It is yet another object of the present invention to provide a system for predicting crash though matching operation between data obtained on driving a vehicle and data learned by neural networks. It is still another object of the present invention to provide a system for evading crash of a vehicle using neural networks to actuate a vehicle safety system for protecting an occupant.

In order to achieve the above-mentioned objects, an apparatus comprising the features of claim 1 and a method comprising the features of claim 9 are provided.

The neural network comprises at least an input layer and an output layer, and the training data are supplied to the input layer while the output layer is supplied with, as teacher data, flags representing expected and unexpected crash, respectively, of the vehicle. In addition, the neural network may comprise a two-dimensional self-organizing competitive learning layer as an intermediate layer.

Other advantages and features of the present invention will be described in detail in the following preferred embodiments thereof.

Brief Description of the Drawings

Fig. 1 is a block diagramm of a conventional system for predicting and evading crash of a vehicle;

Fig. 2 is a schematic view showing a processing element in a typical neural network;

Fig. 3 is a graphical representation of a sigmoid function used as a transfer function for training neural networks;

Fig. 4 is a block diagram of a system for predicting and evading crash of a vehicle using neural networks according to the first embodiment of the present invention;

Fig. 5(a) is a schematic structural diagram of a crash predicting circuit in Fig. 4 realized by a neural network of three layers;

Fig. 5(b) shows an example of an input layer consisting of a two-dimensional array of processing elements of the neural network shown in Fig. 5(a);

Figs. 6(a) and 6(b) are exemplified views picked up, as the training image data supplied to the neural network, at different time instances during driving an experimental vehicle;

Fig. 7 is a view showing an example of an image data obtained during driving a utility vehicle;

Fig. 8 is a view showing another example of an image data obtained during driving a utility vehicle; and

Fig. 9 is a block diagram of a system for predicting and evading crash using neural networks according to the second embodiment of the present invention.

Detailed Description of the Preferred Embodiments

A conventional system for predicting and evading crash of a vehicle is described first to facilitate an understanding of the present invention. Throughout the following detailed description, similar reference numerals refer to similar elements in all figures of the drawing.

In the following description, the term "crash" is used in a wider sense that relates to all unexpected traffic accidents. Accidents other than crash include a turnover or fall of a vehicle, with which the phenomenon of "crash" is associated in some degrees therefore the use of term crash as a cause of traffic accidents.

As shown in Fig. 1, an image pick-up device 21 is mounted at a front portion of an automobile 10 to pick up ever-changing images as analog image data. This image pick-up device 21 is any one of suitable devices such as a charge-coupled-device (CCD) camera. The image data are subject to sampling for a sampling range $_{\Delta}T$ during a predetermined sampling period $_{\Delta}t$. The image data are collected up to crash. In this event, the image pick-up range of the image pick-up device 21 corresponds to a field of view observed through naked eyes.

The image pick-up device 21 is connected to an input interface 22. The analog image data obtained by the image pick-up device 21 are supplied to the input interface 22. The input interface 22 serves as an analog-to-digital converter for converting the analog image data into digital image data. More particularly, the picked up images are digitized by means of dividing the same into tiny pixels (data elements) isolated by grids. It is preferable to eliminate noises and distortions at this stage. The input interface 22 is also connected to a speed sensor 23, a steering gear ratio sensor 24 and a signal processor 30. The speed sensor 23 supplies velocity data to the signal processor 30 through the input interface 22. The velocity data represents an actual velocity of the automobile 10 at the time instant when the image pick-up device 21 picks up an image of a view. Likewise, the steering gear ratio sensor 24 supplies steering gear ratio data to the signal processor 30 through the input interface 22. The steering gear ratio data represents an actual steering gear ratio of the automobile 10.

The signal processor 30 comprises a central processing unit (CPU) 31, a read-only memory (ROM)

IPR2025-00943
Tesla EX1002 Page 1420

32 and a random-access memory (RAM) 33. CPU 31, ROM 32 and RAM 33 are operatively connected to each other through a data bus 34. To evade potentially dangerous objects, CPU 31 carries out calculation operation in response to the image, velocity and steering gear ratio data given through the input interface 22. CPU 31 performs proper functions according to programs stored in ROM 32 and RAM 33. The outputs of the signal processor 30 is transmitted through an output interface 40. ROM 32 stores a table relating to numerical values required for the calculation. It also stores a table representing operational amount for a safety drive ensuring arrangement 50. On the other hand, RAM 33 stores programs for use in calculating an optimum operational amount for the safety drive ensuring arrangement 50. A program for this purpose is disclosed in, for example, Teruo Yatabe, Automation Technique: Intelligent Vehicle, pages 22-28.

The signal processor 30 first determines, according to the picked up image data, whether there is a space available on the roadway to pass through. When there is enough space to pass through and a potentially dangerous object is present on the roadway, the signal processor 30 calculates optimum operational amount for the safety drive ensuring arrangement 50 to operate the same. In Fig. 1, the safety drive ensuring arrangement 50 consists of a steering actuator 51, a throttle actuator 52 and a brake actuator 53. If the signal processor 30 determines that it is necessary to operate these actuators, it produces steering gear ratio command, set velocity command, and brake operation command. The steering actuator 51, the throttle actuator 52 and the brake actuator 53 are operated depending on the condition in response to the steering gear ratio command, the set velocity command and the brake operation command, respectively.

The actuators are for use in actuating occupant protecting mechanism such as a brake device. Operation of these actuators is described now.

The steering actuator 51 is a hydraulic actuator for use in rotating steering wheel (not shown) in an emergency. In this event, the steering wheel is automatically rotated according to the steering gear ratio and rotational direction indicated by the steering gear ratio command. The operational amount of the steering or hydraulic actuator can be controlled in a well-known manner through a servo valve and a hydraulic pump, both of which are not shown in the figure.

The throttle actuator 52 acts to adjust opening amount of a throttle valve (not shown) to decrease speed while evading objects or so on.

The brake actuator 53 performs a function to gradually decrease speed of a vehicle in response to the brake operational command. The brake actuator 53 is also capable of achieving sudden brake operation, if necessary.

As mentioned above, CPU 31 carries out its operation with the tables and programs stored in ROM 32 and RAM 33, respectively, for every one picked up image data. The conventional system is thus disadvantageous in that the calculation operation requires relatively long time interval as mentioned in the preamble of the instant specification.

On the contrary, a system according to the present invention uses image data representing ever-changing views picked up from a vehicle until it suffers from an accident. These image data are used for training a neural network implemented in the present system. After completion of the training, the neural network is implemented in a utility vehicle and serves as a decision making circuit for starting safety driving arrangements to evade crash, which otherwise will certainly be happened. The neural network predicts crash and evades the same by means of properly starting an automatic steering system or a brake system.

A well-known neural network is described first to facilitate an understanding of the present invention and, following which preferred embodiments of the present invention will be described with reference to the drawing.

A neural network is the technological discipline concerned with information processing system, which has been developed and still in their development stage. Such artificial neural network structure is based on our present understanding of biological nervous systems. The artificial neural network is a parallel, distributed information processing structure consisting of processing elements interconnected unidirectional signal channels called connections. Each processing element has a single output connection that branches into as many collateral connections as desired.

A basic function of the processing elements is described below.

As shown in Fig. 2, each processing element can receive any number of incoming functions while it has a single output connection that can be fan out into copies to form multiple output connections. Thus the artificial neural network is by far more simple than the networks in a human brain. Each of the input data x1, x2, ..., xi is multiplied by its corresponding weight coefficient w1, w2, ..., wi, respectively, and the processing element sums the weighted inputs and passes the result through a nonlinearity. Each processing element is characterized by an internal threshold or offset and by the type of nonlinearity and processes a predetermined transfer function to produce an output f(X) corresponding to the sum ($X = \sum xi \cdot wi$). In Fig. 2, xi represents an output of an i-th processing element in an (s-1)-th layer and wi represents a connection strength or the weight from the (s-1)-th layer to the s-th layer. The output f(X) represents energy condition of each processing element. Though the neural networks come in a variety of forms, they can be generally classified into feedforward and recurrent classes. In the latter, the output of each processing element is fed back to other processing elements via weights. As described above, the network has

IPR2025-00943
Tesla EX1002 Page 1421

an energy or an energy function associated with it that will be minimum finally. In other words, the network is considered to have converged and stabilized when outputs no longer change on successive iteration. Means to stabilize the network depends on the algorithm used.

The back propagation neural network is one of the most important and common neural network architecture, which is applied to the present invention. In this embodiment, the neural network is used to determine if there is a possibility of crash. When the neural network detects the possibility of crash, it supplies an operational command to a safety ensuring unit in a manner described below. As well known in the art, the back propagation neural network is a hierarchical design consisting of fully interconnected layers of processing elements. More particularly, the network architecture comprises at least an input layer and an output layer. The network architecture may further comprise additional layer or N hidden layers between the input layer and the output layer where N represents an integer that is equal to or larger than zero. Each layer consists of one or more processing elements that are connected by links with variable weights. The net is trained by initially selecting small random weights and internal thresholds and then presenting all training data repeatedly. Weights are adjusted after every trial using information specifying the correct result until weights converge to an acceptable value. The neural network is thus trained to automatically generate and produce a desired output for an unknown input.

Basic learning operation of the back propagation neural network is as follows. First, input values are supplied to the neural network as the training data to produce output values, each of which is compared with a correct or desired output value (teacher data) to obtain information indicating a difference between the actual and desired outputs. The neural network adjusts the weights to reduce the difference between them. More particularly, the difference can be represented by a well-known mean square error. During training operation, the network adjusts all weights to minimize a cost function equal to the mean square error. Adjustment of the weights is achieved by means of back propagating the error from the output layer to the input layer. This process is continued until the network reaches a satisfactory level of performance. The neural network trained in the above mentioned manner can produce output data based on the input data even for an unknown input pattern.

The generalized delta rule derived with the steepest descent may be used to optimize the learning procedure that involves the presentation of a set of pairs of input and output patterns. The system first uses the input data to produce its own output data and then compares this with the desired output. If there is no difference, no learning takes place and otherwise the weights are changed to reduce the difference. As a result of this it becomes possible to converge the network after a relatively short cycle of training.

To train the net weights on connections are first initialised randomly and input data (training data) are successively supplied to the processing elements in the input layer. Each processing element is fully connected to other processing elements in the next layer where a predetermined calculation operation is carried out. In other words, the training input is fed through to the output. At the output layer the error is found using, for example, a sigmoid function and is propagated back to modify the weight on a connection. The goal is to minimize the error so that the weights are repeatedly adjusted and updated until the network reaches a satisfactory level of performance. A graphical representation of sigmoid functions is shown in Fig. 3.

In this embodiment a sigmoid function as shown in Fig. 3 is applied as the transfer function for the network. The sigmoid function is a bounded differentiable real function that is defined for all real input values and that has a positive derivative everywhere. The central portion of the sigmoid (whether it is near 0 or displaced) is assumed to be roughly linear. With the sigmoid function it becomes possible to establish effective neural network models.

As a sigmoid function parameter in each layer, a y-directional scale and a y-coordinate offset are defined. The y-directional scale is defined for each layer to exhibit exponential variation. This results in improved convergence efficiency of the network.

It is readily understood that other functions may be used as the transfer function. For example, in a sinusoidal function a differential coefficient for the input sum in each processing element is within a range equal to that for the original function. To use the sinusoidal function results in extremely high convergence of training though the hardware for implementing the network may be rather complex in structure.

An embodiment of the present invention is described with reference to Figs. 4 through 9.

Fig. 4 is a block diagram of a system for predicting and evading crash of a vehicle using neural networks according to the first embodiment of the present invention. A system in Fig. 4 is similar in structure and operation to that illustrated in Fig. 1 other than a crash predicting circuit 60. Description of the similar components will thus be omitted by the consideration of evading redundancy. Fig. 5 is a schematic structural diagram of the crash predicting circuit 60 illustrated in Fig. 4 realized by a neural network of three layers.

The crash predicting circuit 60 in this embodiment is implemented by a neural network architecture of a hierarchical design with three layers as shown in Fig. 5(a). The input layer 61 consists of n processing elements 61-1 through 61-n arranged in parallel as a one-dimensional linear form. Each processing element in the input layer 61 is fully connected in series to the processing elements in a hidden layer 62 of the network. The hidden layer 62 is connected to an output

IPR2025-00943
Tesla EX1002 Page 1422

layer 63 of a single processing element to produce an operational command described below. Fig. 5(b) shows an input layer consisting of a two-dimensional array of processing elements. In this event, the image data are supplied to the input layer as a two-dimensional data matrix of n divisions. Basically, the input and the hidden layers can have any geometrical form desired. With the two-dimensional array, the processing elements of each layer may share the same transfer function, and be updated together. At any rate, it should be considered that each processing element is fully interconnected to the other processing elements in the next layer though only a part of which are shown in Fig. 5(a) to evade complexity.

Referring now to Fig. 6 in addition to Fig. 5, illustrated are views picked up, as the image data for use in training the neural network. The image pick-up device 21 picks up ever-changing images as analog image data as described above in conjunction with the conventional system. This image pick-up device 21 is also any one of suitable devices such as a CCD camera. The image pick-up operation is carried out during running of a vehicle at higher speed than a predetermined one. The image data are subject to sampling for a sampling range $_{\Delta}T$ during a predetermined sampling period $_{\Delta}t$. The image data are collected before and just after pseudo crash. The image pick-up range of the image pick-up device 21 corresponds to a field of view observed through naked eyes. A view shown in Fig. 6(a) is picked up when a station wagon (estate car) 80a on the opposite lane comes across the center line. A view shown in Fig. 6(b) is picked up when an automobile 80b suddenly appears from a blind corner of a cross-street. These ever-changing images are collected as the training data for the neural network.

The image data effectively used for the crash evasive purpose are those which allow continuous recognition of the ever-changing views before and just after pseudo crash. With this respect, the image pick-up device 21 picks up the images of a vehicle or other obstructions located at a relatively short headway. In addition, the picked up images preferably are distinct reflections of the outside views.

The data elements consisting of one image are simultaneously supplied to the input layer 61 in parallel. In other words, each data element is supplied to the respective processing element of the input layer 61. The digital image data may be normalized before being supplied to the input layer 61 to increase a data processing speed. However, each processing element of the input layer 61 essentially receives the data element obtained by dividing the image data previously. The data elements are subjected to feature extraction when supplied to the hidden layer 62.

In typical image processing, feature extraction is carried out according to any one of various methods of pattern recognition to clearly identify shapes, forms or configurations of images. The feature-extracted data are quantized for facilitate subsequent calculations. In this event, separate analytical procedure is used for region partitioning or for extraction of configuration strokes. In other words, a particular program is necessary for each unit operation such as region partitioning, feature extraction, vectorization and so on. Compared with this, the prediction system according to the present invention requires no program based on each operation or procedure because a unique algorithm is established on completion of network training. This single algorithm allows to perform necessary functions without using separate algorithms or programs.

In a preferred embodiment, the feature extraction is directed to the configuration of an object defining the driving lanes such as shoulders, curbs, guard rails or the center line. The feature may also be extracted on regions such as carriageways. The neural network learns these configurations and regions during training process. This process is continued until the network reaches a satisfactory level of performance. The neural network is thus trained while carrying out feature extraction on the input image. Weights are adjusted after every trial on the quantized image data, so that the latest training data is weighted according to the latest result of adjustment and then supplied to the hidden layer 62. In addition, the neural network can be trained with image data including an object at time-varying positions. In this event, any one of suitable methods may be used for digital image processing.

In the present embodiment, each digital data indicative of ever-changing view at a certain sampling time instance is divided into n data elements. A product of n represents a positive integer which is equal in number to the processing elements in the input layer 61. In other words, the series of time sequential data is picked up as continuous n data elements to be supplied in parallel to the n by m processing elements in the input layer 61 as the training data. At the same time, an operational signal is supplied to the output layer 63 of the network as teacher data. The operational signal may be a logic "1" for representing crash of the automobile 10 after elapse of a predetermined time interval from the sampling time instant corresponding to the image data just having been supplied to the input layer 61.

In the same manner, the picked up image data and its corresponding teacher data are successively supplied to the crash predicting circuit 60. The crash predicting circuit 60 is continuously trained until the network reaches a satisfactory level of performance. After completion of training, the network is capable of matching the picked up image with the possibility of crash. The accuracy of prediction is improved by means of supplying images for a case of "safe" state to the neural network on learning.

The neural network thus learns the relative position between the vehicle on which it is mounted and objects at a short headway. As a result of this learning, the crash predicting circuit 60 enables to prediction of crash

IPR2025-00943
Tesla EX1002 Page 1423

expected to be happened a few seconds later according to this relative position. While outside views change every moment and a vehicle in practice encounters various objects and situations, a series of repeated training can yield stereotyped data patterns.

The neural network program that has already been trained can be memorized in a read only memory (ROM) as an application. In this event the network program is memorized after being compiled and translated into a machine language. The ROM is implemented in a predetermined IC chip or the like as an inherent circuit. The IC chip is mounted on a circuit for the air bag system in an automobile.

As mentioned above, the crash predicting circuit 60 supplies the operational signal to the safety drive ensuring arrangement 50 when it predicts occurrence of crash. In response to this operational signal the safety drive ensuring arrangement 50 can perform proper function to evade crash.

For more clear understanding of the present invention, two cases are described those results in "safe" state of the automobile 80c, 80d. Fig. 7 shows an exemplified image including an oncoming vehicle 80c running on the opposite lane. The situation being far from danger as shown in Fig. 7 may allow the system of the present invention to bypass the crash predicting circuit 60. Alternatively, the crash predicting circuit 60 may produce an operational signal of logic "0" to represent this "safe" condition.

A view shown in Fig. 8 represents a situation when a vehicle 80d on the opposite lane comes across the center line in the far distance ahead. The vehicle 80d is going to return to the lane where it ought to be. The subsequent image data indicate that the oncoming vehicle 80d takes an action to evade crash. In other words, the oncoming vehicle 80d is expected to return to the proper lane before the vehicle mounting the crash predicting circuit 60 passes by the vehicle 80d. Accordingly, the crash predicting circuit 60 determines that there are no hazardous objects ahead.

If a vehicle on the opposite lane comes across the center line or a vehicle suddenly appears from a blind corner of a cross-street as shown in Figs. 5(a) and 5(b), the crash predicting circuit 60 carries out prediction operation in accordance with the image data showing these situations. Expected hazards make the crash predicting circuit 60 actuate the safety drive ensuring arrangement 50 in the manner described above.

Another embodiment of the present invention will be described below in which the neural network comprises an intermediate layer having a self-organization function and a competitive learning function to positively respond to various unknown data with less training data. As well known in the art, in the self-organization a network modifies itself in response to inputs. Examples of the use of self-organizing training include the competitive learning law applied to the present embodiment.

As shown in Fig. 9 the neural network according to this embodiment comprises a two-dimensional self-organized competitive learning layer 64 interposed between the input layer 61 and the hidden layer 62. The two-dimensional self-organized competitive learning layer 64 is referred as to the two-dimensional Kohonen layers (2D-K layer) which in this embodiment comprises p by q layers consisting of a two-dimensional array of processing elements. The input layer 61 may consists of either one or two-dimensional array of processing elements. The 2D-K layer 64 can have any geometrical form desired. In this embodiment, it is also considered that each processing element is fully interconnected to the other processing elements in the next layer though only a part of which are shown in Fig. 9 to evade complexity.

The processing elements in the 2D-K layer 64 compete with one another to determine the "winner" on the basis of minimum distance. More particularly, a predetermined distance can be obtained by, in this embodiment, n processing elements for each set of the input data. The similarity for each of the n input data corresponds to the distance to select similar combination of processing elements. The selected processing elements becomes "winner" for facilitating determination on attributes of unknown data.

More particularly, the winning three Kohonen's processing elements are determined among the fourth processing elements to supply output data. Unknown data are preprocessed on the basis of classification for the input data due to the self-organization on learning. The output value thereof is supplied to the subsequent hidden layer.

With an additional normalization layer 65 may be interposed between the input layer 61 and the 2D-K layer 64 as shown in Fig. 9. With this normalization layer 65, the learning efficiency in the 2D-K layer 64 will be sufficiently improved. Addition of the 2D-K layer 64 contributes to a surprising number of information processing capabilities for unknown data as well as a remarkably improved convergence efficiency on learning.

The neural network having the 2D-K layer can be completed by means of expanding the above mentioned back propagation method so that the learning procedure can be determined in a similar manner as in the back propagation method.

The self-organization requires that the system uses, during adaptation of initial several thousands times, no other information other than the incoming patterns and no data are fed back from the output layer. After completion of self-organization the network is trained according to the back propagation algorithm. The neural network having a structure according to this embodiment can be trained with less data for a shorter period of training cycle.

In the above mentioned second embodiment, the neural network already trained can be coded by using a programming language such as C-language. The net-

IPR2025-00943
Tesla EX1002 Page 1424

work may be used as an imperative application system or packaged as a control microprocessor. In this event, the network can be memorized in a read only memory for every one type of commercial vehicles.

For the portion of algorithm that is established readily in logical, a well-known expert system may be applied to achieve a prediction system using a combination of logic circuit for the neural network and the expert system.

While the above embodiments have thus been described in conjunction with automatic crash evasive operation, it is possible to give precedence to the driver's operation. For example, it is possible to issue appropriate warnings to the driver before actuation of the safety drive ensuring arrangement 50. For this purpose, an audible signal such as an alarm sound may be generated to alert the driver to potential hazards.

Claims

1. An apparatus for predicting and evading potential crash of a vehicle and predicting real crash of a vehicle and acting correspondingly, comprising:

   image pick-up means mounted on the vehicle for picking up images of actual ever-changing views when the vehicle is on running to produce actual image data;
   crash predicting means comprising a neural network associated with said image pick-up means, said crash predicting means successively supplied with the actual image data for predicting potential occurence of crash between the vehicle and potentially dangerous objects to produce an operational signal when there is possibility of crash,
   safety drive ensuring means connected to said crash predicting means for actuating, in response to the operational signal, predicting potential or real crash,
   characterized in that
   an occupant protecting mechanism including an airbag is provided,
   said neural network being adapted and trained to predict real occurence of a crash and being connected with said airbag, to release said airbag in case of real crash,
   a safety drive ensuring arrangement is provided,
   said neural network being adapted and trained to predict occurence of potential crash and being connected with said safety drive ensuring arrangement to perform proper function to evade crash,
   said neural network therefore being trained with training data to predict the possibility of potential occurence and real occurence of crash, the training data representing ever-changing views previously picked up during driving of the vehicle and after real crash.

2. Apparatus according to claim 1, **characterized** in that

   means for producing an audible signal to alert the driver to potential occurence of crash are provided.

3. Apparatus according to claim 1, **characterized** in that

   said neural network comprises at least an input layer and an output layer, and
   training data are supplied to the input layer while the output layer is supplied with, as teacher data,
   flags representing "crash expected" and "no crash exprected", respectively, of the vehicle.

4. Apparatus according to claim 3, **characterized** in that

   said neural network comprises a two-dimensional self-organizing competitive learning layer as an intermediate layer.

5. Apparatus according to claim 1, **characterized** in that

   said neural network is coded after completion of learning and implemented in the vehicle.

6. Apparatus according to claim 1, **characterized** in that

   said safetey drive ensuring means includes a steering actuator and the occupant protecting mechanism includes a steering system of the vehicle.

7. Apparatus according to claim 1, **characterized** in that

   said safety drive ensuring means includes a throttle actuator and the occupant protecting mechanism includes a throttle system of the vehicle.

8. Apparatus according to claim 1, **characterized** in that

   said safety drive ensuring means includes a brake actuator and the occupant protecting mechanism includes a brake system of the vehicle.

IPR2025-00943
Tesla EX1002 Page 1425

9. Method for predicting and evading potential crash of a vehicle and for predicting real crash of a vehicle and acting correspondingly, comprising:

image pick-up means mounted on the vehicle for picking up images of actual ever-changing views when the vehicle is on running to produce actual image data;
crash predicting means comprising a neural network associated with said image pick-up means, said crash predicting means successively supplied with the actual image data for predicting potential occurence of crash between the vehicle and potentially dangerous objects to produce an operational signal when there is possibility of crash,
safety drive ensuring means, connected to said crash predicting means for actuating, in response to the operational signal, the occupant protecting mechanism including an airbag,
and a safety drive ensuring arrangement to perform proper function to evade crash, which are operatively connected thereto and equipped in the vehicle,
whereby
said neural network contains previously taken image data formed of successive scenes for causing accidents and is trained by a back propagation method for realizing conditions of causing said accidents, whereby the training of said neural network is performed by supplying training data to an input layer of said neural network while an output layer of said neural network is supplied with, as teacher data, flags representing "crash expected" and "no crash expected",
respectively, of the vehicle,
said neural network evaluates the actual image data by itself,
said neural network judges if the vehicle collides with the objects noticed in the actual image data of the image pick-up means, and outputs said operational signal in case of prediction of occurence of a crash with said object.

10. Method according to claim 9, **characterized** in that

said neural network comprises a two-dimensional self-organizing competitive learning layer as an intermediate layer.

11. Method according to claim 9, **characterized** in that

said neural network is coded after completion of learning and implemented in the vehicle.

Patentansprüche

1. Vorrichtung zum Vorhersagen und Vermeiden eines potentiellen Zusammenstoßes eines Fahrzeugs und zum Vorhersagen eines realen Zusammenstoßes eines Fahrzeugs und zum entsprechenden Agieren, mit:

einer Bildaufnahmeeinrichtung, die auf dem Fahrzeug angebracht ist, um Bilder von aktuellen, ständig wechselnden Ansichten aufzunehmen, wenn das Fahrzeug fährt, um aktuelle Bilddaten zu erzeugen;
einer Zusammenstoß-Vorhersageeinrichtung mit einem neuronalen Netzwerk, das der Bildaufnahmeeinrichtung zugeordnet ist, wobei der Zusammenstoß-Vorhersageeinrichtung sukzessive die aktuellen Bilddaten zum Vorhersagen eines potentiellen Auftretens eines Zusammenstoßes zwischen dem Fahrzeug und potentiell gefährlichen Objekten zugeführt sind, um ein Operationssignal zu erzeugen, wenn die Möglichkeit eines Zusammenstoßes besteht,
einer Einrichtung zur Gewährleistung der Fahrsicherheit, die mit der Zusammenstoß-Vorhersageeinrichtung verbunden ist, um in Reaktion auf das Operationssignal, das einen potentiellen oder realen Zusammenstoß vorhersagt, betätigt zu werden,
dadurch **gekennzeichnet,** daß
ein Insassenschutzmechanismus mit einem Airbag vorgesehen ist,
das neuronale Netzwerk zur Vorhersage eines realen Auftretens eines Zusammenstoßes angepaßt und trainiert ist und mit dem Airbag verbunden ist, um den Airbag im Fall eines realen Zusammenstoßes auszulösen,
eine Anordnung zur Gewährleistung der Fahrsicherheit vorgesehen ist,
das neuronale Netzwerk zur Vorhersage des Auftretens eines potentiellen Zusammenstoßes angepaßt und trainiert ist und mit der Anordnung zur Gewährleistung der Fahrsicherheit verbunden ist, um eine geeignete Funktion zur Vermeidung eines Zusammenstoßes durchzuführen,
das neuronale Netzwerk daher mit Trainingsdaten trainiert ist, um die Möglichkeit des potentiellen Auftretens und des realen Auftretens eines Zusammenstoßes vorherzusagen, wobei die Trainingsdaten ständig wechselnde, früher aufgenommene Ansichten darstellen, die während des Fahrens eines Fahrzeuges und nach einem realen Zusammenstoß aufgenommen sind.

2. Vorrichtung nach Anspruch 1,

IPR2025-00943
Tesla EX1002 Page 1426

dadurch **gekennzeichnet,**

daß eine Einrichtung zur Erzeugung eines hörbaren Signals vorgesehen ist, um den Fahrer vor dem potentiellen Auftreten eines Zusammenstoßes zu warnen.

3. Vorrichtung nach Anspruch 1, dadurch **gekennzeichnet,**

daß das neuronale Netzwerk mindestens eine Eingabeschicht und eine Ausgabeschicht aufweist, und Trainingsdaten der Eingabeschicht zugeführt werden, während der Ausgabeschicht als Lerndaten Kennzeichen zugeführt werden, die "Zusammenstoß erwartet" beziehungsweise "kein Zusammenstoß erwartet" des Fahrzeugs repräsentieren.

4. Vorrichtung nach Anspruch 3, dadurch **gekennzeichnet,**

daß das neuronale Netzwerk eine zweidimensionale, selbstorganisierende kompetitive Lernschicht als Zwischenschicht aufweist.

5. Vorrichtung nach Anspruch 1, dadurch **gekennzeichnet,**

daß das neuronale Netzwerk nach Abschluß des Lernens kodiert und in dem Fahrzeug implementiert ist.

6. Vorrichtung nach Anspruch 1, dadurch **gekennzeichnet,**

daß die Einrichtung zur Gewährleistung der Fahrsicherheit eine Steuerungsbetätigung aufweist und der Insassenschutzmechanismus ein Steuerungssystem des Fahrzeugs aufweist.

7. Vorrichtung nach Anspruch 1, dadurch **gekennzeichnet,**

daß die Einrichtung zur Gewährleistung der Fahrsicherheit eine Drosselklappenbetätigung aufweist und der Insassenschutzmechanismus ein Drosselklappensystem des Fahrzeugs aufweist.

8. Vorrichtung nach Anspruch 1, dadurch **gekennzeichnet,**

daß die Einrichtung zur Gewährleistung der Fahrsicherheit eine Bremsenbetätigung aufweist und der Insassenschutzmechanismus ein Bremsensystem des Fahrzeugs aufweist.

9. Verfahren zum Vorhersagen und Vermeiden eines potentiellen Zusammenstoßes eines Fahrzeugs und zum Vorhersagen eines realen Zusammenstoßes eines Fahrzeugs und zum entsprechenden Agieren, mit:

einer Bildaufnahmeeinrichtung, die auf dem Fahrzeug zur Aufnahme von Bildern von aktuellen, ständig wechselnden Ansichten angebracht ist, wenn das Fahrzeug fährt, um aktuelle Bilddaten zu erzeugen;
einer Zusammenstoß-Vorhersageeinrichtung mit einem neuronalen Netzwerk, das der Bildaufnahmeeinrichtung zugeordnet ist, wobei der Zusammenstoß-Vorhersageeinrichtung sukzessive die aktuellen Bilddaten zur Vorhersage des potentiellen Auftretens eines Zusammenstoßes zwischen dem Fahrzeug und potentiell gefährlichen Objekten zugeführt werden, um ein Operationssignal zu erzeugen, wenn die Möglichkeit eines Zusammenstoßes besteht,
einer Einrichtung zur Gewährleistung der Fahrsicherheit, die mit der Zusammenstoß-Vorhersageeinrichtung verbunden ist, um in Reaktion auf ein Operationssignal betätigt zu werden, wobei der Insassenschutzmechanismus einen Airbag aufweist, und
einer Anordnung zur Gewährleistung der Fahrsicherheit, um eine geeignete Funktion zum Vermeiden eines Zusammenstoßes durchzuführen, die operativ damit verbunden und in dem Fahrzeug vorgesehen ist, wobei das neuronale Netzwerk früher aufgenommene Bilddaten enthält, die von aufeinanderfolgenden Szenen beim Verursachen eines Unfalls gebildet werden und durch ein Rück-Propagation-Verfahren zum Realisieren von Verhältnissen, die Unfälle verursachen, trainiert ist, wobei das Training des neuronalen Netzwerks durch Zuführen von Trainingsdaten zu einer Eingabeschicht des neuronalen Netzwerks durchgeführt wird, während eine Ausgabeschicht des neuronalen Netzwerks mit Kennzeichen als Lerndaten versorgt wird, die "Zusammenstoß erwartet" und "kein Zusammenstoß erwartet" des Fahrzeugs repräsentieren,
das neuronale Netzwerk die aktuellen Bilddaten selbst bewertet,
das neuronale Netzwerk beurteilt, ob das Fahrzeug mit den beobachteten Objekten in den aktuellen Bilddaten der Bildaufnahmeeinrichtung kollidiert und das Operationssignal im Falle einer Vorhersage des Auftretens eines Zusammenstoßes mit dem Objekt ausgibt.

10. Verfahren nach Anspruch 9, dadurch **gekennzeichnet,**

IPR2025-00943
Tesla EX1002 Page 1427

daß das neuronale Netzwerk eine zweidimensionale, selbstorganisierende kompetitive Lernschicht als Zwischenschicht aufweist.

11. Verfahren nach Anspruch 9, dadurch **gekennzeichnet,**

daß das neuronale Netzwerk nach Vollendung des Lernens kodiert und in dem Fahrzeug implementiert wird.

Revendications

1. Appareil pour prédire et prévenir une collision potentielle d'un véhicule et pour prédire une collision réelle d'un véhicule et agir en conséquence, comprenant:

un moyen de saisie d'image monté sur le véhicule pour saisir, quand le véhicule roule, des images de vues réelles constamment changeantes, afin de produire des données d'images réelles;
un moyen de prédiction de collision comprenant un réseau neuronal associé audit moyen de saisie d'image, ledit moyen de prédiction de collision étant alimenté successivement avec les données d'images pour prédire la survenue potentielle d'une collision entre le véhicule et des objets potentiellement dangereux afin de produire un signal d'opération lorsqu'il existe un risque de collision;
un moyen assurant une conduite sûre, relié au moyen de prédiction de collision pour un actionnement, en réponse au signal d'opération, et pour une prédiction de collision potentielle ou réelle,
caractérisé en ce que:
un mécanisme de protection d'occupant comprenant un coussin gonflable est présent,
ledit réseau neuronal étant adapté et éduqué de manière à prédire la survenue réelle d'une collision et étant relié audit coussin gonflable pour déclencher le fonctionnement dudit coussin gonflable en cas de collision réelle,
un agencement assurant une conduite sûre est prévu, ledit réseau neuronal étant adapté et éduqué de manière à prédire la survenue d'une collision potentielle et étant relié audit agencement assurant une conduite sûre afin de remplir une fonction appropriée pour prévenir une collision,
ledit réseau neuronal étant par conséquent éduqué avec des données d'éducation de manière à prédire le risque de la survenue potentielle et de la survenue réelle d'une collision, les données d'éducation représentant des vues constamment changeantes saisies antérieurement pendant que le véhicule roule et après une collision réelle.

2. Appareil selon la revendication 1, caractérisé en ce qu'un moyen servant à produire un signal audible pour avertir le conducteur d'une survenue potentielle de collision est présent.

3. Appareil selon la revendication 1, caractérisé en ce que ledit réseau neuronal comprend au moins une couche d'entrée et une couche de sortie, et que les données d'éducation sont fournies à la couche d'entrée pendant que la couche de sortie reçoit, en tant que données d'éducateur, des indicateurs représentant une "collision prévue" et "aucune collision prévue" du véhicule, respectivement.

4. Appareil selon la revendication 3, caractérisé en ce que le réseau neuronal comprend, en tant que couche intermédiaire, une couche bidimensionnelle d'apprentissage compétitrif s'auto-organisant

5. Appareil selon la revendication 1, caractérisé en ce que ledit réseau neuronal est codé après achèvement de l'enseignement et est intégré dans le véhicule.

6. Appareil selon la revendication 1, caractérisé en ce que ledit moyen assurant une conduite sûre comprend un actionneur de direction et le mécanisme de protection d'occupant comprend le système de direction du véhicule.

7. Appareil selon la revendication 1, caractérisé en ce que ledit moyen assurant une conduite sûre comprend un actionneur de papillon des gaz et le mécanisme de protection d'occupant comprend le système de papillon des gaz du véhicule.

8. Appareil selon la revendication 1, caractérisé en ce que le moyen assurant une conduite sûre comprend un actionneur de freins et le mécanisme de protection d'occupant comprend le système de freins du véhicule.

9. Procédé pour prédire et prévenir une collision portentielle d'un véhicule et pour prédire une collision réelle d'un véhicule et agir en conséquence, comprenant:

un moyen de saisie d'image monté sur le véhicule pour saisir, quand le véhicule roule, des images de vues réelles constamment changeantes afin de produire des données d'images réelles;
un moyen de prédiction de collision comprenant un réseau neuronal associé audit moyen de saisie d'images, ledit moyen de prédiction

IPR2025-00943
Tesla EX1002 Page 1428

de collision recevant de façon successive les données d'images réelles pour prédire la survenue potentielle d'une collision entre le véhicule et des objets potentiellement dangereux afin de produire un signal d'opération quand il existe un risque de collision,

un moyen assurant une conduite sûre, relié audit moyen de prédiction de collision pour actionner, en réponse au signal d'opération, le mécanisme de protection comprenant un coussin gonflable,et un agencement asurant une conduite sûre pour remplir une fonction appropriée afin de prévenir une collision, qui y sont reliés fonctionnellement et montés dans le véhicule,

grâce à quoi ledit réseau neuronal contient des données d'images saisies antérieurement, formées de scènes successives entraînant des accidents et est éduqué par un procédé de rétro-propagation servant à réaliser les conditions entraînant lesdits accidents, l'éducation dudit réseau neuronal étant effectuée par fourniture de données d'éducation à une couche d'entrée dudit réseau neuronal pendant qu'une couche de sortie dudit réseau neuronal reçoit, en tant que données d'éducateur, des indicateurs représentant une "collision prévue" et "aucune collision prévue" du véhicule, respectivement,

ledit réseau neuronal évalue par lui-même les données d'images réelles,

ledit réseau neuronal juge si le véhicule heurte les objets mentionnés dans les données d'images réelles du moyen de saisie d'images et émet ledit signal d'opération dans le cas d'une prédiction de survenue d'une collision avec ledit objet.

10. Procédé selon la revendication 9, caractérisé en ce que ledit réseau neuronal comprend, en tant que couche intermédiaire, une couche bidimensionnelle d'apprentissage compétitif s'auto-organisant.

11. Procédé selon la revendication 9, caractérisé en ce que ledit réseau neuronal est codé après achèvement de l'enseignement et est intégré dans le véhicule.

IPR2025-00943
Tesla EX1002 Page 1429

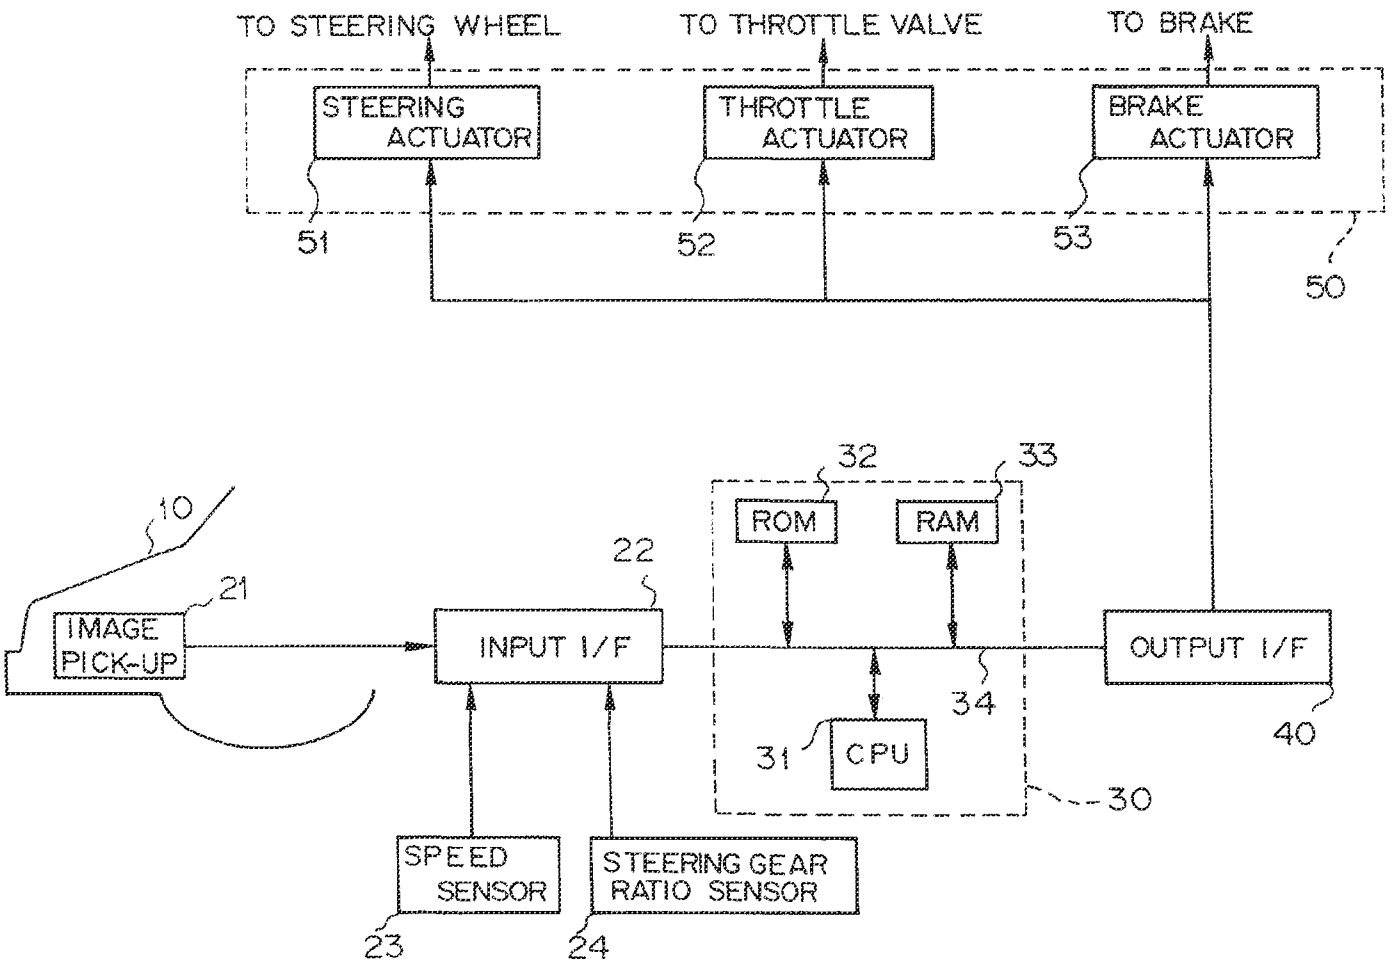


FIG. 1 PRIOR ART

EP 0 582 236 B1

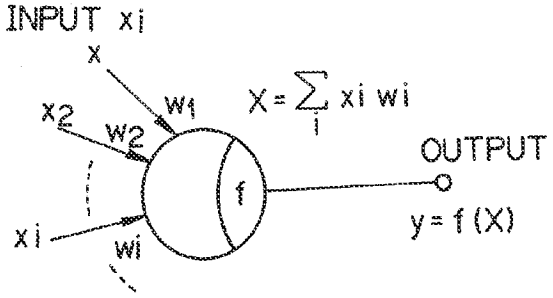


FIG. 2
PRIOR ART

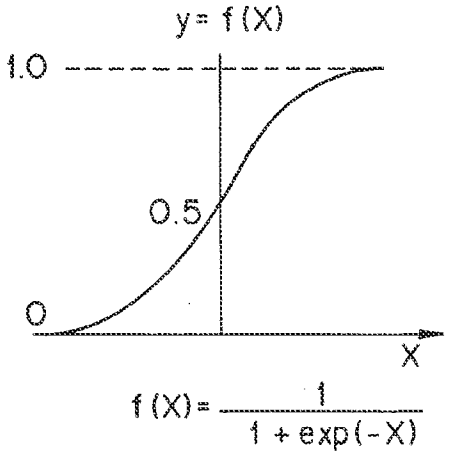


FIG. 3
PRIOR ART

IPR2025-00943
Tesla EX1002 Page 1431

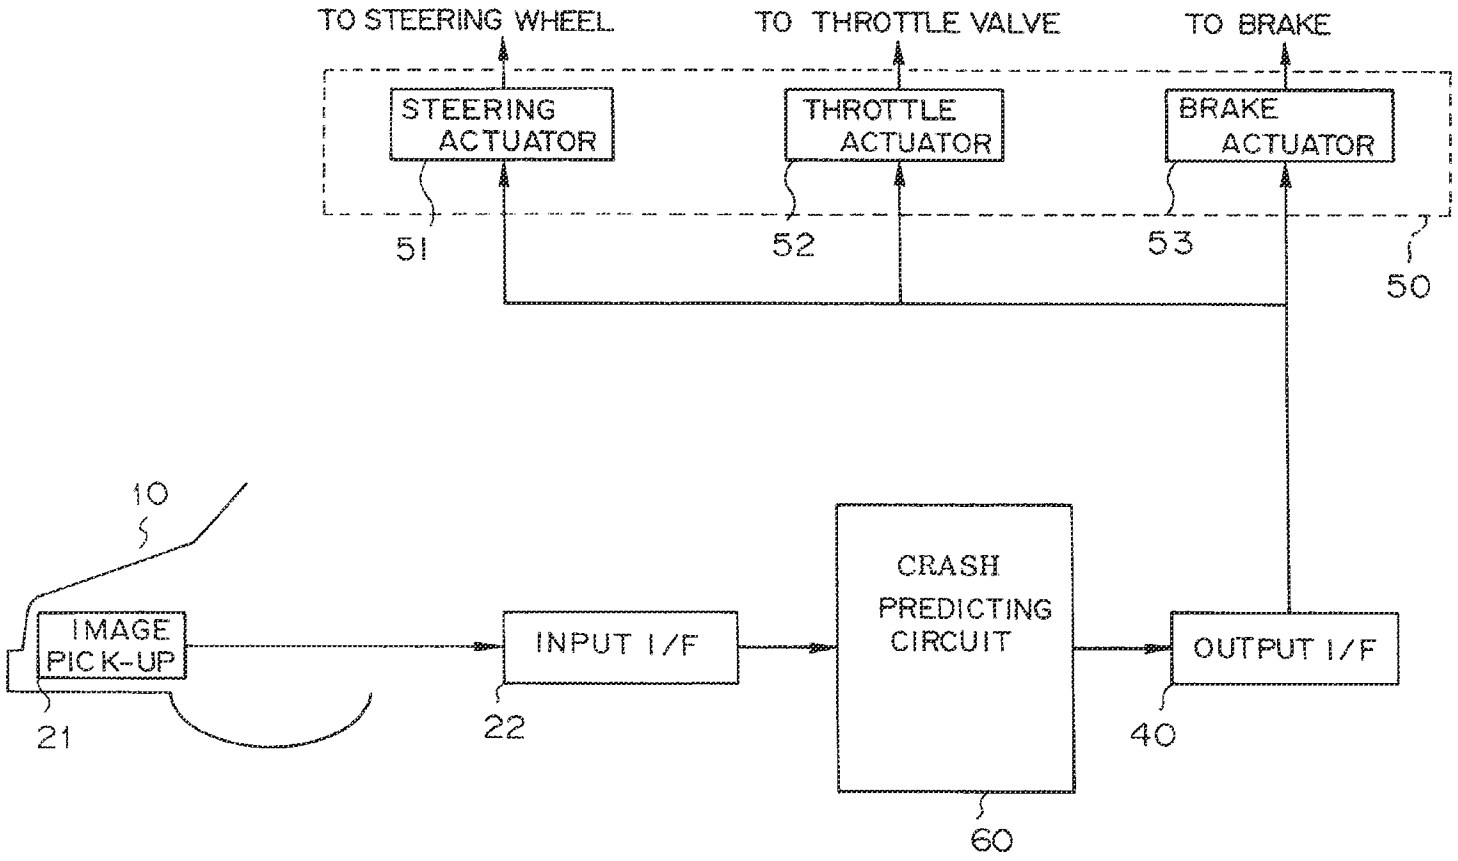
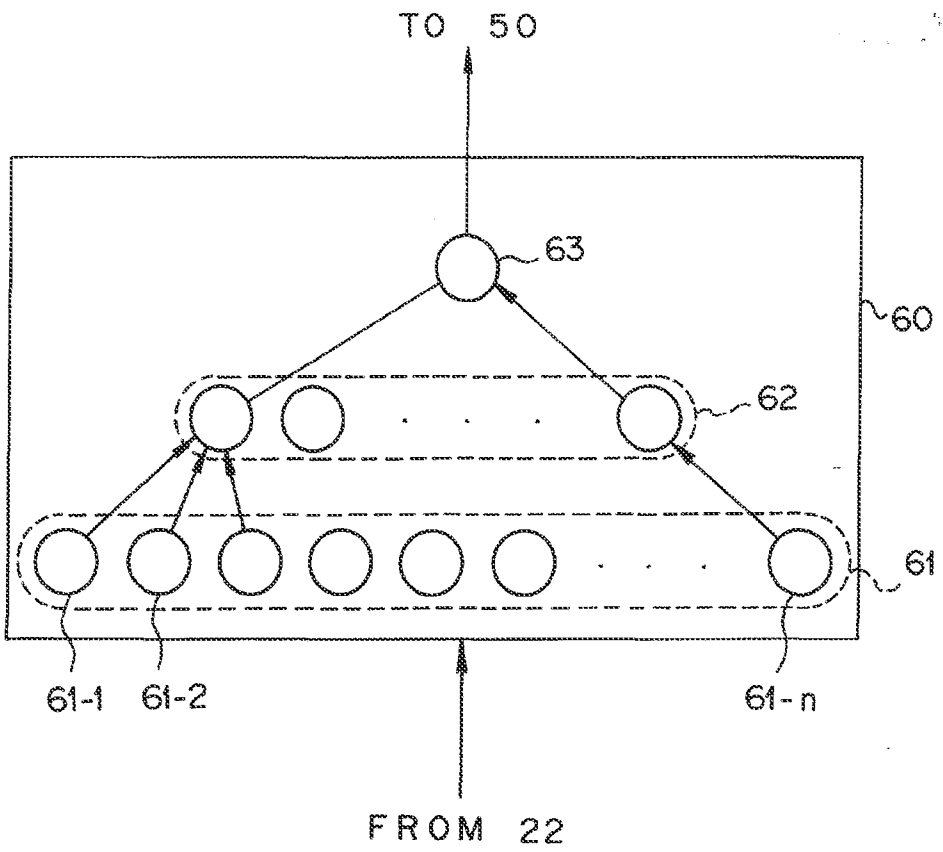


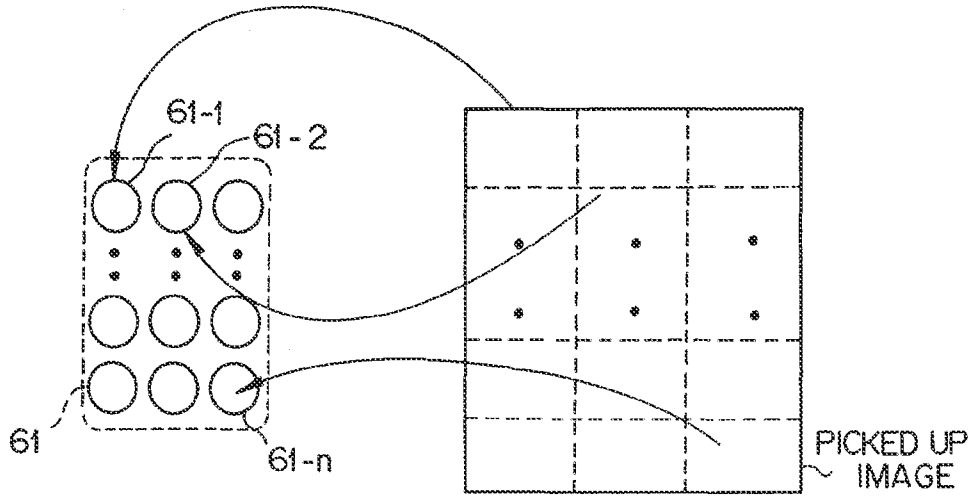
FIG. 4

EP 0 582 236 B1

IPR2025-00943
Tesla EX1002 Page 1432



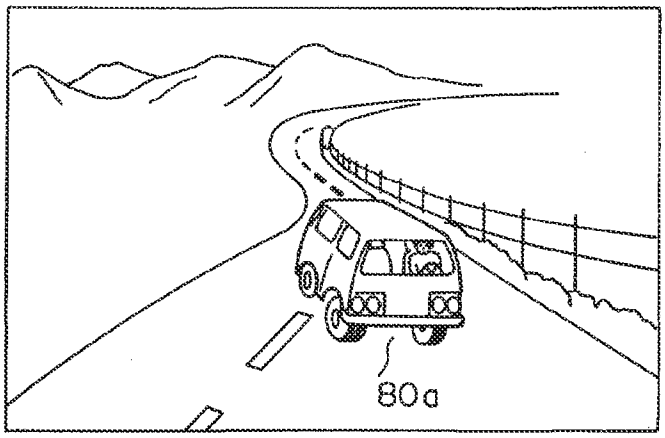
(a)



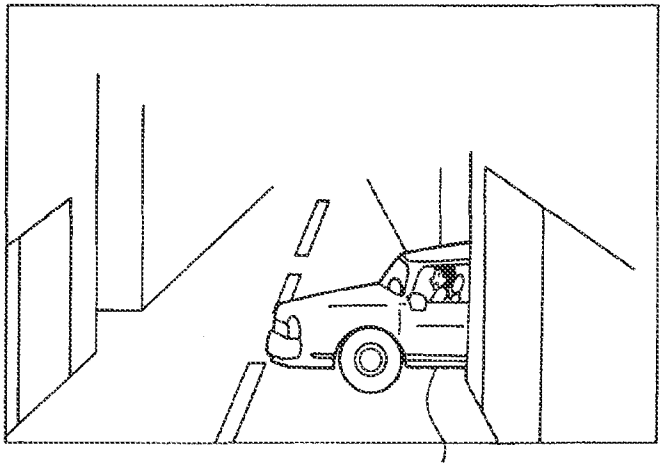
(b)

FIG. 5

IPR2025-00943
Tesla EX1002 Page 1433



(a)



(b)

FIG. 6

IPR2025-00943
Tesla EX1002 Page 1434

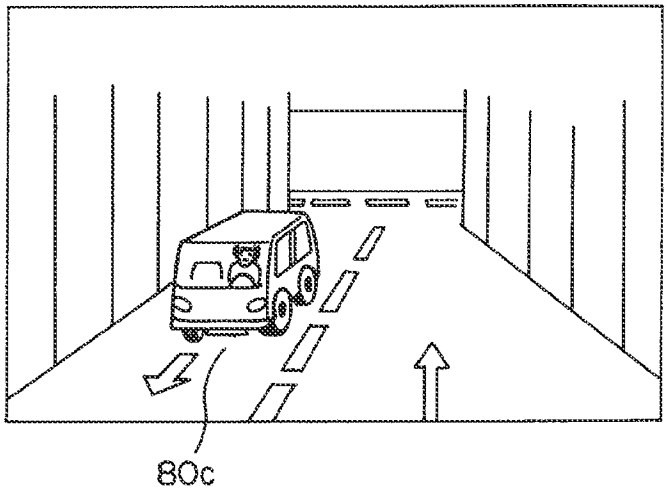


FIG. 7

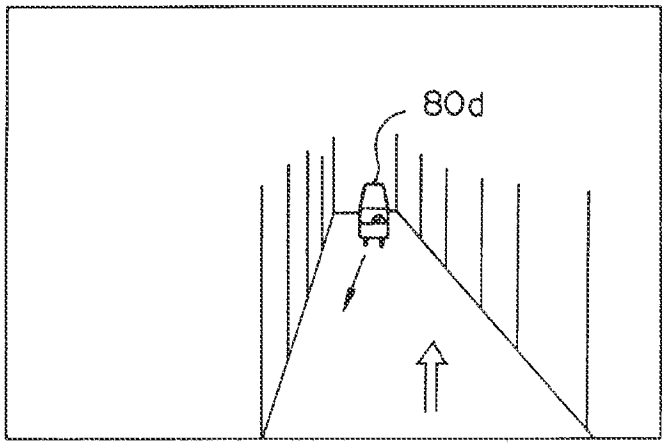


FIG. 8

IPR2025-00943
Tesla EX1002 Page 1435

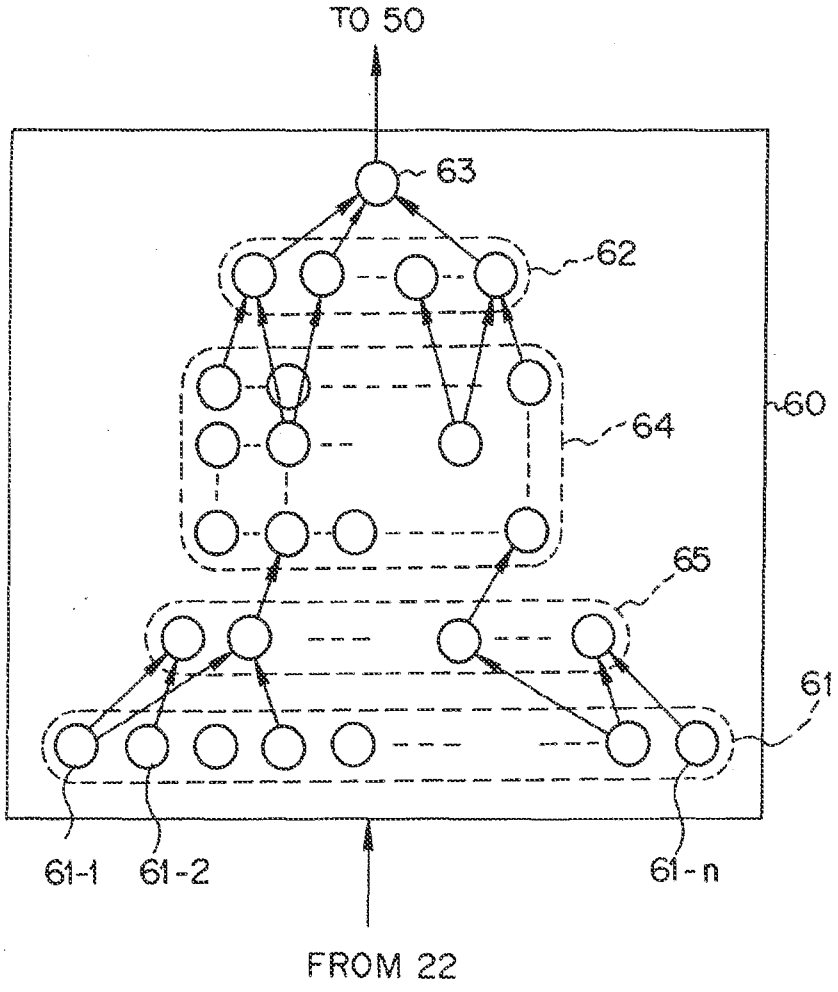


FIG. 9

IPR2025-00943
Tesla EX1002 Page 1436

Espacenet

Bibliographic data: EP1236126 (A1) — 2002-09-04

SYSTEM FOR DETECTING OBSTACLES TO VEHICLE MOTION

Inventor(s): STEIN GIDEON [IL]; SHASHUA AMNON [IL] ± (STEIN, GIDEON, ; SHASHUA, AMNON)

Applicant(s): MOBILEYE INC [US] ± (MOBILEYE, INC, ; MOBILEYE TECHNOLOGIES, LTD)

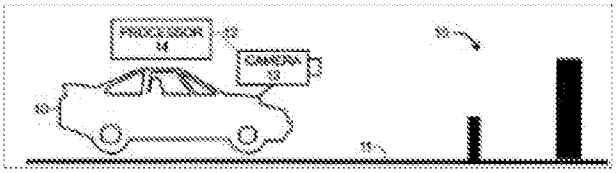
Classification: - international: ***B60R21/00; B60R21/01; G01S11/12; G06F17/10; G06F17/11; G08G1/16;*** *B60R21/0134; G01S5/16;* (IPC1-7): G08G1/04; G08G1/16
- cooperative: **B60R21/013 (EP); G01S11/12 (EP);** B60R21/0134 (EP); G01S5/163 (EP)

Application number: EP20000982211 20001127 Global Dossier

Priority number (s): WO2000US32144 20001127 ; US19990167587P 19991126

Also published as: EP1236126 (A4) EP1236126 (B1) AT307370 (T) AU1927101 (A) CA2392578 (A1) more

Abstract not available for EP1236126 (A1)
Abstract of corresponding document: WO0139018 (A1)



A time-to-contact estimate determination system is disclosed for generating an estimate as to the time-to-contact of a vehicle (10) moving along a roadway (11) with an obstacle (15). The time-to-contact estimate determination system comprises an image receiver (camera 13) and a processor (14). The image receiver is configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway. The processor is configured to process the image information received by the image receiver to generate a time-to-contact of the vehicle with the obstacle.

IPR2025-00943
Tesla EX1002 Page 1437

Espacenet

Description: EP1236126 (A1) — 2002-09-04

SYSTEM FOR DETECTING OBSTACLES TO VEHICLE MOTION

Description of EP1236126 (B1)

A high quality text as facsimile in your desired language may be available amongst the following family members:

CA2392578 (A1) DE60023339 (T2) ES2250218 (T3) WO0139018 (A1)

The EPO does not accept any responsibility for the accuracy of data and information originating from other authorities than the EPO; in particular, the EPO does not guarantee that they are complete, up-to-date or fit for specific purposes.

[0001] The invention relates generally to the field of systems and methods for estimating time-to-contact between a moving vehicle and an obstacle and more specifically to systems and methods that estimate time-to-contact using successively-recorded images recorded along the vehicle's path of motion.

BACKGROUND OF THE INVENTION

[0002] Accurate estimation of the time-to-contact between a vehicle and obstacles is an important component in autonomous driving and computer vision-based driving assistance. Using computer vision techniques to provide assistance while driving, instead of mechanical sensors, allows for the use of the information that is recorded for use in estimating vehicle movement to also be used in estimating ego-motion identifying lanes and the like, without the need for calibration between sensors as would be necessary with mechanical sensors. This can reduce the cost of the arrangements provided to provide time-to-contact estimates and maintenance that may be required therefor.

[0003] US Patent 5,646,612 describes apparatus for avoiding collision of a vehicle with an object. The apparatus uses a laser radar system to determine location of the object and a speed sensor to determine speed of the vehicle. Position and speed data provided by the radar and the speed sensor are used to determine an estimate for a time to collision of the vehicle and the object.

SUMMARY OF THE INVENTION

[0004] The invention provides a new and improved system that estimates time-to-

contact of a vehicle with an obstacle using successively-recorded images recorded along the vehicle's path of motion.

[0005] In accordance with the present invention, a time-to-contact estimate determination system for generating an estimate as to the time-to-contact of a vehicle moving along a roadway with an obstacle comprises:
A. an image receiver configured to receive image information relating to a series of at least two images as the vehicle moves along the roadway; and
B. a processor characterized in that the processor is configured to determine a scaling factor that defines a ratio between dimensions of the obstacle in the images and uses the ratio to generate a time-to-contact estimate of the vehicle with the obstacle.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:
FIG. 1 schematically depicts a vehicle moving on a roadway and including a time-to-contact estimation system constructed in accordance with the invention;
FIGS. 2A and 2B are useful in understanding the apparent horizontal motion of an obstacle in a series of images as a function of horizontal position of the obstacle relative to the vehicle, which is useful in estimating the likelihood of the vehicle contacting the obstacle; and
FIG. 3 is a flow chart depicting operations performed by the time-to-contact estimation system in estimating time-to-contact of the vehicle with an obstacle.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

[0007] FIG. 1 schematically depicts a vehicle 10 moving on a roadway 11 and including a time-to-contact estimation system 12 constructed in accordance with the invention. The vehicle 10 may be any kind of vehicle 10 that may move on the roadway 11, including, but not limited to automobiles, trucks, buses and the like. The time-to-contact estimation system 12 includes a camera 13 and a processor 14. The camera 13 is mounted on the vehicle 10 and is preferably pointed in a forward direction, that is, in the direction in which the vehicle would normally move, to record successive images as the vehicle moves over the roadway. Preferably as the camera 13 records each image, it will provide the image to the processor 14. The processor 14, in turn, will process information that it obtains from the successive images, possibly along with other information, such as information from the vehicle's speedometer (not separately shown) to estimate a time-to-contact value corresponding to the estimated time-to-contact, if any, of the vehicle 10 with one or more obstacles, generally identified by reference numeral 15. The processor 14 may also be mounted in or on the vehicle 11 and may form part thereof. The time-to-contact estimates generated by the processor 14 may be used for a number of things, including, but not limited to autonomous driving by the vehicle, providing assistance in collision avoidance, and the like. Operations performed by the processor 14 in estimating time-to-contact will be described in connection with the flow chart depicted in FIG. 3.

[0008] Before proceeding further, it would be helpful to provide some background for the operations performed by the collision detection and time-to-contact estimation processor 14 depicted in FIG. 1 in detecting possible obstacles to motion of the vehicle

IPR2025-00943
Tesla EX1002 Page 1439

10 and the time period to contact between the vehicle 10 and the respective obstacles, if any. Generally, the operations performed by the processor 14 can be divided into two phases, namely, an obstacle detection phase and a time to contact estimation phase. It will be appreciated that these phases may overlap in connection with various ones of the obstacles, with the processor 14 engaging in the obstacle detection phase to attempt to detect new obstacles while it is engaging in the time to contact estimation phase to determine time-to-contact with previously-detected obstacles.

[0009] During the obstacle detection phase, two general operations are performed. Initially, using a roadway deletion operator, as between successive images Ψ, Ψ',..., the portions of the images that comprise projections of the roadway are identified. The portions of the images that are identified as comprising projections of the roadway will be ignored. Using an obstacle detection operator, portions of the images other than those that comprise projections of the roadway are analyzed to identify obstacles. After the obstacles are identified, during the time to contact estimation phase, a time-to-contact estimate is generated for the identified obstacles. The time-to-contact estimate indicates whether the vehicle 10 and obstacle are moving closer together, farther apart, or maintaining a constant separation. In particular, if the time-to-contact estimate for a particular obstacle is positive, the vehicle 10 and the obstacle are moving closer together, if the estimate is infinite, the separation is constant, and if the estimate is negative, the vehicle 10 and the obstacle are moving farther apart.

[0010] It will be appreciated that the time-to-contact estimate generated for a particular pair of images Ψ and Ψ' will provide an estimate as to the time-to-contact at the time that the later image Ψ' was recorded. However, the time-to-contact estimate generated as described herein does not take into account other information that may be useful in determining the likelihood of whether the vehicle 10 will contact the obstacle if their relative motion remains constant. For example, ego-motion information (which may be provided as generated, for example, using a methodology described in the Stein patent application) and information as to the shape of the roadway and expected path of the vehicle may be useful in determining whether the vehicle will likely contact the obstacle. If, for example, the obstacle is in the same lane as the vehicle 10, and if the time-to-contact estimate is positive, indicating that the separation of the vehicle 10 and obstacle is decreasing, it may be determined that the vehicle 10 will likely contact the obstacle if the vehicle's motion is not changed. In addition, using the same information, strategies may be developed for avoiding contact, such as changing speed, changing lanes or other strategies as will be apparent to those skilled in the art.

[0011] As noted above, during the obstacle detection phase, a road detection operator and an obstacle detection operator are used to facilitate identification of obstacles. Although the obstacle detection operator alone could be used in identification ofobstacles, since the roadway is not an obstacle, and since typically the projection of the roadway typically comprises relatively large portions of the respective images Ψ and Ψ', those portions can be ignored in applying the obstacle detection operator. The roadway detection operator initially tessellates, into corresponding regions "R," the image Ψ and an image Ψ", where image Ψ" is the warp of image Ψ' toward image Ψ using the estimated motion of the vehicle between the time at which image Ψ was recorded to the time at which image Ψ' is recorded, and generates values "Q" as follows:
Image available on "Original document"
The estimated motion that is used to generate the warped image Ψ" comprises the translation and rotation of the vehicle 10 as between the point in time at which image Ψ

IPR2025-00943
Tesla EX1002 Page 1440

was recorded and the point in time at which image Ψ' was recorded, and may be an initial guess based on, for example, the vehicle's speed as provided by a speedometer, or the estimated motion as generated as described in the Stein patent application. It will be appreciated that the warped image Ψ' generally reflects an estimate of what the image would have been at the time that the image Ψ was recorded. This holds for regions that are projections of the roadway, but not necessarily for other regions, and particular will not hold for regions that have a vertical extent, which will be the case for obstacles. Accordingly, regions in images Ψ and Ψ' for which the value "Q" is below a selected threshold will be deemed to be regions that are projections of the roadway, and other regions will be deemed not to be regions that are projections of the roadway. Similarly, regions in image Ψ' which were warped to regions in image Ψ' that are deemed to be regions that are projections of the roadway, are also deemed regions that are projections of the roadway, and other regions in image Ψ' are deemed not to be projections of the roadway.

[0012] After the roadway detection operator has been used to identify regions in images Ψ and Ψ' that are not projections of the roadway, those regions are processed using the obstacle detection operator to detect obstacles in those regions. Generally, it will be appreciated that, unlike the roadway or elements, artifacts such as roadway markings, shadows cast on the roadway surface, and the like, all of which have only a horizontal extent, obstacles have a vertical extent as well as a horizontal extent. Accordingly, as between two successive images Ψ and Ψ', artifacts such as markings on the roadway surface, shadows cast on the roadway surface, and the like, will exhibit expansion along the images' horizontal axes and little or no expansion along the vertical axes. On the other hand, obstacles, which have vertical extent, will exhibit more uniform expansion along the images' vertical axes as well as along the horizontal axes.

[0013] With this observation, in the obstacle detection phase, the images are processed using two filters, namely, a roadway filter and an obstacle filter. The roadway filter filters out portions of the roadway and associated artifacts on the images Ψ and Ψ', and leaves the obstacles. The affine motion of a patch as between image Ψ and image Ψ' that is a projection of an obstacle exhibits a generally uniform scaling as around a "focus of expansion" ("FOE") for the obstacle. That is, if the FOE of an obstacle is at pointpFOE(xo,y0)in the image Ψ, the motion vector(u, v) of a projection of a pointP(X,Y,Z)on the obstacle, which projects onto pointp(x,y)in the image Ψ, will have components u = s (x - x 0)
image available on "Original document"
v = s (y - y 0)Image available on "Original document"
where "s" is a constant that reflects the uniform scaling. The components of the motion vector indicate the horizontal and vertical motion of a projection of a pointP(X,Y,Z)in three dimensional space as between the projectionp(x,y)in image Ψ and the projectionp'(x',y')in image Ψ'. If the coordinates of the FOE are not known, the motion vector will still reflect a uniform scaling, and will also reflect a translation u = sx + x f
image available on "Original document"
v = sy + y fImage available on "Original document"
where xf= -sx0and yf= -sy0. It will be appreciated that, in equations (4) and (5), the "sx" and "sy" terms reflect the uniform scaling, and the xfand yfterms reflect the translation. Since, for the motion vector, the components u=x'-x and v=y'-y, where "x" and "y" are the coordinates of the projection of a point in three-dimensional space in image Ψ and "x'" and "y'" are the coordinates of the projection of the same point in three-dimensional space in image Ψ', and so
image available on "Original document"

Tesla EX1002 Page 1441 IPR2025-00943

Substituting equation (6) into the constant-brightness criterion ul x + vl y + l t = 0Image available on "Original document"
yieldsImage available on "Original document"
which can be writtenImage available on "Original document"
where the superscript "T" represents the transpose operation. Regions of images Ψ, Ψ' for which equation (9) hold represent obstacles.

[0014] If one or more obstacles is detected during the obstacle detection phase, the processor 14 can generate an estimate of the time-to-contact, if any, of the vehicle 10 with the obstacle. Although the time-to-contact will reflect the distance between the vehicle and the obstacle, it will be appreciated that in many circumstances the time-to-contact is a more useful metric. The distance of the obstacle to the image plane of the camera 13, which, in turn, reflects the distance of the obstacle to the vehicle 10, at the point in time T-ΔT at which image Ψ was recorded is Z+ΔZ, and the distance at the point in time T at which image Ψ' was recorded is Z. (It should be noted that the values of both "ΔT" and "ΔZ" are positive). If there is no change in relative motion between the vehicle 10 and the obstacle, the distance Z to the obstacle will be closed in a time period T, that is, the time-to-contact is T.

[0015] The processor 14 estimates the time-to-contact T in relation to the ratio of the scaling of the obstacle as recorded in the images Ψ and Ψ', if any, and specifically in relation to the scaling of the vertical dimension. The vertical dimension of the obstacle in image Ψ is y 2 = fY Z + Δ Z
Image available on "Original document"
where "Y" refers to the height of the obstacle in three-dimensional space and "f" is the focal length of the camera 13. Similarly, the vertical dimension of the obstacle in image Ψ' is y 1 = fY ZImage available on "Original document"
since the height "Y" of the obstacle in three-dimensional space does not change as between the times at which the images Ψ and Ψ' are recorded. The scaling factor, or ratio, of the vertical dimensions "S" in the images is S = y 1 y 2 = Z + Δ Z Z = 1 + Δ Z ZImage available on "Original document"
In addition, assuming that there is no change in motion as between the vehicle 10 and the obstacle, the relative speed Δ Z Δ TImage available on "Original document"
between the vehicle 10 and the obstacle during the time period between the points in time at which the images Ψ and Ψ' were recorded, will be the same as the relative speed Z TImage available on "Original document"
during the time period between the point in time at which the image Ψ' was recorded and the point in time at which the vehicle 10 would contact the obstacle. That is Δ Z Δ T = Z TImage available on "Original document"
Rearranging equation 13, T Δ T = Z Δ ZImage available on "Original document"
Combining equation (14) and equation (12), S = 1 + Δ Z Z = 1 + Δ T TImage available on "Original document"
Rearranging equation (15) to solve for T, the time-to-contact, T = 1 S - 1 Δ TImage available on "Original document"
It will be appreciated that, if the value of ΔZ is positive and non-zero, from equation (12), the value of S will be greater than one, in which case the value of the factor 1 S - 1Image available on "Original document"
in equation (16) will be greater than zero. In that case, the time-to-collision T will be positive. It will be appreciated that this can occur if the distance separating the vehicle 10 and the obstacle is decreasing. On the other hand, if the value of the ratio S is equal to one, which can occur if the vertical dimension of the obstacle in the image Ψ is the same as the vertical dimension in the image Ψ', the value of the factor 1 S - 1Image

IPR2025-00943
Tesla EX1002 Page 1442

available on "Original document"
in equation (16) is infinite, in which case the time-to-contact T will also be infinite. It will be appreciated that, when that occurs, if the vehicle 10 is moving, the obstacle will also be moving, and it will be moving at the same speed as the vehicle 10, and in that case, the separation between the vehicle 10 and the obstacle will be constant. Finally, if the value of the ratio S is less than one, which can occur if the vertical dimension of the obstacle in the image Ψ is greater than the vertical dimension in the image Ψ', the value of the factor 1 S - 1Image available on "Original document"
in equation (16) is negative, in which case the time-to-contact T will also be negative. It will be appreciated that, when that occurs, the obstacle will also be moving, and it will be moving at a speed that is greater than, and in a direction away from, the vehicle 10. It will be appreciated that, using equation (16), the processor 14 can estimate the time-to-contact using only information from the images Ψ and Ψ' and the time period ΔT, without needing any information as to the actual distance between the vehicle 10 and the obstacle.

[0016] As noted above, the time-to-contact value generated as described above actually reflects the rate at which the separation between the vehicle 10 and the obstacle is decreasing at the point in time that the image Ψ' is recorded. Depending on the particular positions and velocities (which reflect direction as well as speed) of the vehicle 10 and the obstacle at any point in time, as well as the sizes, and primarily the horizontal dimensions, of the vehicle 10 and the obstacle, the vehicle and obstacle may or may not actually come into contact. This will be described in connection with FIGs. 2A and 2B. With reference to FIG. 2A, that FIG. schematically depicts a portion of a curved roadway 20 with three lanes, including a center lane represented by line 21C, a lane to the left of the center lane represented by line 21L and a lane to the right of the center lane represented by line 21R. The center of curvature of the roadway is at O. The vehicle 10 is traveling in the center lane 21C, and obstacles, which comprise other vehicles, are traveling in all three lanes. To simplify the explanation, the coordinate system will be deemed to be moving with the vehicle 10. It will also be assumed that a series of images is recorded, including a first image Ψ when the obstacles are at locations a, a', a", a second image Ψ" when the obstacles are at locations b, b', b", a third image Ψ"' when the obstacles are at locations c, c', c", a fourth image Ψ"" when the obstacles are at locations d, d', d", and so forth. In this example, the separations between vehicle 10 and all three obstacles are decreasing. Graphs of the separations and the horizontal coordinates of the respective obstacles in the series of images (it will be appreciated that the vertical coordinates of the respective obstacles in the images will not change) are depicted in FIG. 2B, with the separation with the obstacle in lane 21C being represented by solid circles, the separation with the obstacle in lane 21 L being represented by hollow circles, and the separation with the obstacle in lane 21R being represented by crosses ("+"). In FIG. 2B, the vehicle 10 will be considered to be located at the origin, that is, where the "x-coordinate" is zero and the Z-coordinate is zero, although it will be appreciated that, because the vehicle 10 will have a non-zero horizontal dimension, the portion of the horizontal dimension of the width subtended by the vehicle extends to the left and right of the origin by an amount associated with the width of the vehicle. As shown in FIG. 2B, all three obstacles will initially approach the vehicle 10, but at some point the separations between vehicle 10 and the obstacles in the left and right lanes 21 R and 21L will again increase. However, the separation between obstacle in the same lane 21C as that of the vehicle 10 will continue to decrease. If the time intervals between times at which the images are recorded arc uniform the horizontal coordinate of the projection of the obstacle in lane 21C will uniformly approach zero. For obstacles in the other lanes 21L and 21R, the movement

IPR2025-00943
Tesla EX1002 Page 1443

of the horizontal coordinate will not be uniform. Using the information as to the uniformity of the progression of the horizontal coordinates of the projection of the vehicles across the successive images, the likelihood as to whether the vehicle will contact an obstacle can be determined.

[0017] With this background, operations performed by the processor 14 in connection with determining the time-to-contact will be described in connection with the flow chart depicted in FIG. 3. With reference to FIG. 3, after receiving a new image Ψ', the processor 14 initially estimates the ego-motion of the vehicle 10 (step 300) using, for example, a methodology described in connection with the Stein patent application. Using the estimated ego-motion and lane information, if available, the processor 14 extrapolates the future path of the vehicle (step 301) and identifies a "danger zone" along the extrapolated future path (step 302). In performing step 302, the processor 14 can make use of geometric assumptions that are based on the camera's calibration parameters, and the danger zone can, for example, comprise a trapezoidal region along the image of the roadway. The processor 14 then applies the roadway detection operator (step 303) and obstacle detection operator (step 304) to those portions to identify obstacles that are in or near the extrapolated path. Thereafter, if the processor 14 will determine whether there were any obstacles identified after application of the obstacle detection operator (step 305), and, if not return to step 300 to receive the next image.

[0018] On the other hand, if the processor 14 determines in step 305 that one or more obstacles have been identified in step 304, the processor 14 will examine the next few images Ψ", Ψ"', Ψ"",... to determine whether the images contain projections of the respective obstacle(s) (step 306). Tracking obstacle(s) whose projections have been detected in one or two images through successive images provides verification that the obstacles in fact exist and are not artifacts in the respective images. For any obstacles whose projections are recorded in a predetermined number of subsequent images, the processor 14 will determine that the obstacle(s) have been verified (step 307), and, for any verified obstacles, the processor 14 will determine the value of the scaling factor "S" for the obstacle as between successive images (step 308). After determining the value of the scaling factor, the processor 14 uses that value and the value ΔT in equation (16) to determine the time-to-contact T (step 309). Thereafter, the processor can make use of the time-to-contact value in connection with, for example, providing collision avoidance assistance to the vehicle or driver, if, for example, the time-to-contact value is positive. As noted above, if the time-to-contact value is not positive, there will be no contact between the vehicle 10 and the obstacle.

[0019] The invention provides a number of advantages. In particular, the invention provides an arrangement for estimating the time-to-contact of a vehicle 10 with an obstacle directly from successively-recorded images of the obstacle, without requiring other mechanical or electronic sensors, which can be expensive to install and maintain.

[0020] It will be appreciated that a number of changes and modifications may be made to the time-to-contact estimation system 12 described above. For example, instead of using the scaling of the vertical dimension of obstacles as between image Ψ and Ψ', the system 12 may use the scaling of the horizontal dimension, or both.

[0021] In addition, it will be appreciated that images Ψ and Ψ' are preferably rectified as, for example, described in the Stein patent application, so that their image planes are perpendicular to the roadway plane.

IPR2025-00943
Tesla EX1002 Page 1444

[0022] Furthermore, it will be appreciated that roadway detection filters and obstacle detection filters can be used other than or in addition to those described herein. For example, a filter such as one described in the aforementioned Stein patent application can be used to identify patches that contain projections of obstacles.

[0023] It will be appreciated that a system in accordance with the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof, any portion of which may be controlled by a suitable program. Any program may in whole or in part comprise part of or be stored on the system in a conventional manner, or it may in whole or in part be provided in to the system over a network or other mechanism for transferring information in a conventional manner. In addition, it will be appreciated that the system may be operated and/or otherwise controlled by means of information provided by an operator using operator input elements (not shown) which may be connected directly to the system or which may transfer the information to the system over a network or other mechanism for transferring information in a conventional manner.

Espacenet

Claims: EP1236126 (A1) — 2002-09-04

SYSTEM FOR DETECTING OBSTACLES TO VEHICLE MOTION

Claims of EP1236126 (B1)

A high quality text as facsimile in your desired language may be available amongst the following family members:

CA2392578 (A1) ES2250218 (T3) WO0139018 (A1)

- Original claims
- Claims tree

The EPO does not accept any responsibility for the accuracy of data and information originating from other authorities than the EPO; in particular, the EPO does not guarantee that they are complete, up-to-date or fit for specific purposes.

1. A time-to-contact estimate determination system (12) for generating an estimate as to the time-to-contact of a vehicle (10) moving along a roadway (11) with an obstacle (15) comprising:
A. an image receiver (13) configured to receive image information relating to a series of at least two images as the vehicle moves along the roadway; and
B. a processor (14) ; characterized in that the processor is configured to determine a scaling factor that defines a ratio between dimensions of the obstacle in the images and uses the ratio to generate a time-to-contact estimate of the vehicle with the obstacle.

2. A system according to claim 1, wherein the scaling factor defines a ratio between vertical dimensions of the obstacle in the images and the processor uses the ratio to estimate the time-to-contact.

3. A system according to claim 1, wherein the scaling factor defines a ratio between horizontal dimensions of the obstacle in the images and the processor uses the ratio to estimate the time-to-contact.

4. A system according to any of the preceding claims, wherein the at least two images comprise more than two images.

5. A system according to claim 4, wherein the processor (14) processes the image information to determine a lateral displacement of the object relative to a position of the vehicle.

6. A according to claim 5, wherein the processor (14) determines a likelihood of collision responsive to whether or not the lateral displacement substantially uniformly approaches zero.

IPR2025-00943
Tesla EX1002 Page 1447

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Veröffentlichungsnummer:
(11) Publication number: **EP 1 236 126 A0**
(11) Numéro de publication:

Internationale Anmeldung veröffentlicht durch die Weltorganisation für geistiges Eigentum unter der Nummer:

WO 01/039018 (art. 158 des EPÜ).

International application published by the World Intellectual Property Organisation under number:

WO 01/039018 (art. 158 of the EPC).

Demande internationale publiée par l'Organisation Mondiale de la Propriété sous le numéro:

WO 01/039018 (art. 158 de la CBE).

Espacenet

Bibliographic data: EP1257971 (A2) — 2002-11-20

SYSTEM AND METHOD FOR ESTIMATING EGO-MOTION OF A MOVING VEHICLE USING SUCCESSIVE IMAGES RECORDED ALONG THE VEHICLE'S PATH OF MOTION

Inventor(s): STEIN GIDEON [IL]; SHASHUA AMNON [IL]; MANO OFER [IL] ± (STEIN, GIDEON, ; SHASHUA, AMNON, ; MANO, OFER)

Applicant(s): MOBILEYE INC [US] ± (MOBILEYE, INC)

Classification: **- international:** ***G01S11/12; G05D1/02; G06T1/00; G06T7/20; G08G1/16;*** *G01S5/16;* (IPC1-7): G05D1/02; G06T7/20

- cooperative: **G01S11/12 (EP);** G01S5/163 (EP)

Application number: EP20000980706 20001127 Global Dossier

Priority number (s): WO2000US32143 20001127 ; US19990167587P 19991126 ; US20000230166P 20000901

Also published as: EP1257971 (A4) AU1793301 (A) CA2392652 (A1) JP2003515827 (A) WO0139120 (A2) more

Abstract not available for EP1257971 (A2)
Abstract of corresponding document: WO0139120 (A2)



An ego-motion determination system is disclosed for generating an estimate as to the ego-motion of a vehicle moving along a roadway. The ego-motion determination system includes an image information receiver and a processor. The image information receiver is configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway. The processor is configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle, including the translation of the vehicle in the forward direction and the rotation of the vehicle around a vertical axis as between, for example, successive images.;An ego-motion determination system (fig. 1, item 12) is disclosed for generating an estimate as to the ego-motion of a vehicle (fig. 1, item 10) moving along a roadway (fig. 1, item 11). The ego-motion determination system includes an image information receiver (fig. 1, item 13) and a processor (fig. 1, item 14). The image information receiver (fig. 1, item 13) is configured to receive image information relating

to a series of at least two images recorded as the vehicle moves along a roadway (fig. 1, item 11). The processor is configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle, including the translation of the vehicle in the forward direction and the rotation of the vehicle around a vertical axis as between, for example, successive images (see fig. 3)

IPR2025-00943
Tesla EX1002 Page 1450

Espacenet

Description: EP1257971 (A2) — 2002-11-20

SYSTEM AND METHOD FOR ESTIMATING EGO-MOTION OF A MOVING VEHICLE USING SUCCESSIVE IMAGES RECORDED ALONG THE VEHICLE'S PATH OF MOTION

Description not available for EP1257971 (A2)
Description of corresponding document: WO0139120 (A2)

A high quality text as facsimile in your desired language may be available amongst the following family members:

CA2392652 (A1) WO0139120 (A2)

The EPO does not accept any responsibility for the accuracy of data and information originating from other authorities than the EPO; in particular, the EPO does not guarantee that they are complete, up-to-date or fit for specific purposes.

SYSTEM AND METHOD FOR ESTIMATING EGO-MOTION OF A MOVING VEHICLE USING
SUCCESSIVE IMAGES RECORDE ALONG THE VEHICLE'S PATH OF MOTION
FIELD OF THE INVENTION
The invention relates generally to the field of systems and methods for estimating egomotion (that is,"self-motion) of a moving vehicle, and more specifically to systems and methods that estimate ego-motion using successively-recorded images recorded along the vehicle's path of motion.

BACKGROUND OF THE INVENTION
Accurate estimation of the ego- ("self-") motion of a vehicle relative to a roadway is an important component in autonomous driving and computer vision-based driving assistance. Using computer vision techniques to provide assistance while driving, instead of mechanical sensors, allows for the use of the information that is recorded for use in estimating vehicle movement to also be used in detecting obstacles, identifying lanes and the like, without the need for calibration between sensors as would be necessary with mechanical sensors. This reduces cost and maintenance.

There are several problems in estimating ego-motion of a vehicle. Typically, roads have few feature points, if any. The most obvious features in a road, such as lane markings, have a generally linear structure, whereas background image structures, such as those associated with other vehicles, buildings, trees, and the like, will typically have many feature points. This will make image-or optical-flow-based estimation difficult in

practice. In addition, typically images that are recorded for ego-motion estimation will contain a large amount of"outlier"information that is either not useful in estimating ego-motion, or that may result in poor estimation. For example, in estimating of ego-motion relative to a fixed structure such as a road, images of objects that are moving relative to the road, such as other moving vehicles and even moving wipers, will contribute false information for the ego-motion estimation. In addition, conditions that degrade image quality, such as raindrops and glare, will also make accurate ego-motion estimation difficult.

SUMMARY OF THE INVENTION

The invention provides a new and improved system and method for estimating ego-motion using successively-recorded images recorded along the vehicle's path of motion..

In brief summary, the invention provides an ego-motion determination system for generating an estimate as to the ego-motion of a vehicle moving along a roadway. The ego-motion determination system includes an image information receiver and a processor. The image information receiver is configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway. The processor is configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle, including the translation of the vehicle in the forward direction and the rotation of the vehicle around a vertical axis as between, for example, successive images.

Several methodologies are disclosed for generating the ego-motion estimate of the vehicle.

BRIEF DESCRIPTION OF THE DRAWINGS

This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which :
FIG. 1 schematically depicts a vehicle moving on a roadway and including an ego-motion estimation system constructed in accordance with the invention ;
FIG. 2 is a flow chart depicting operations performed by the ego-motion estimation system in determining ego-motion of the vehicle in accordance with one methodology ; and
FIG. 3 is a flow chart depicting operations performed by the ego-motion estimation system in determining ego-motion of the vehicle in accordance with a second methodology.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

FIG. 1 schematically depicts a vehicle 10 moving on a roadway 11 and including an egomotion estimation system 12 constructed in accordance with the invention. The vehicle 10 may be any kind of vehicle 10 that may move on the roadway 11, including, but not limited to automobiles, trucks, buses and the like. The ego-motion estimation system 12 includes a camera 13 and a egomotion estimation system processor 14. The camera 13 is mounted on the vehicle 10 and is preferably pointed in a forward direction, that is, in the direction in which the vehicle would normally move, to record successive images as the vehicle moves over the roadway. Preferably as the camera 13 records each image, it will provide the image to the ego-motion estimation system processor 14. The ego-motion estimation system processor 14, in turn, will process information that it obtains from the successive images, possibly along with other

IPR2025-00943
Tesla EX1002 Page 1452

information, such as information from the vehicle's speedometer (not separately shown) to determine the ego-motion (that is, the selfmotion) of the vehicle relative to the roadway 11. The ego-motion estimation system processor 14 may also be mounted in or on the vehicle 11 and may form part thereof. The ego-motion estimates generated by the ego-motion estimation system processor 14 may be used for a number of things, including, but not limited to obstacle and lane detection, autonomous driving by the vehicle, perhaps also using positioning information from, for example, the global positioning system ("GPS") and roadway mapping information from a number of sources known to those skilled in the art, and the like. Operations performed by the ego-motion estimation system processor 14 in determining egomotion of the vehicle 10 will be described in connection with the flow charts depicted in FIGS. 2 and 3.

Before proceeding further, it would be helpful to provide background of the operations performed by the ego-motion estimation system processor 14 depicted in FIG. 1. Generally, as between two images T and T'the ego-motion estimation system processor 14 attempts to determine the translation t = (tx, ty, tz) and rotation w = (WX, wy, wz) (where"T"refers to the transpose operation, t, refers to translation along the respective"X,""Y"and"Z"axes, and w, refers to rotation around the respective axis) of the camera 13 affixed to the vehicle 10. Since the camera 13 is affixed to the vehicle 10, the translation and rotation of the camera 13 will also conform to the translation and rotation of the vehicle 10. In that case, for a point p = (x,y)T (where "x" and "y" are coordinates of apoint orfeature inthe image) thatis aprojection of apoint P = (X, Y, Z) (where "X,""Y"and"Z"are coordinates of the point in three-dimensional space), the flow vector for the point, that is, the vector indicating the motion of the same point in three dimensional space from is position in the image T to the image S", has components (u, v)
image available on "Original document"

where"X"in equations (1) refers to the matrix cross-product and
Image available on "Original document"

where"f"is the focal length of the camera 13, which is presumed to be known. The roadway on which the vehicle 10 is traveling is modeled as a plane. The equation for points on a plane is AX+ BY+ BZ = 1 (3), where the"X"and"Y"axes correspond to horizontal"x"and vertical"y"axes of the image plane, and the"Z"axis is perpendicular to the image plane. The camera 13 may be tilted slightly downwardly to increase the amount of the image that corresponds to the road, and reduce the amount of the image that corresponds to other features, such as other traffic, buildings, trees and the like, and in that case, the images T and T'will be rectified so that the images'z axis will be parallel to the plane of the roadway, as will be described below.

Dividing equation (3) by "Z" provides
1/Z = ax + by + c (4), where a=Alf, b=Blf and c=C. Substituting equation (4) into equations (1) results in
image available on "Original document"

Expanding equations (5) results in
image available on "Original document"

Equations (6) and (7) are a special case (the"calibrated camera 13"case) of an eight

parameter model for a camera 13 moving relative to a plane : U = alx+ a2y+ a3+ a7x 2 + a8xy (8) v-a4x+ asy+ a6 + axy+ agyz (9)
Given the flow vector (u, v), one can recover the parameters a,, i=1,..., 8, from which one can recover the motion parameters t and w.

One problem arises in connection with the methodology described above in connection with equations (1) through (9), namely, given the large number of parameters ai whose values need to be determined, it is difficult to devise a satisfactory method that will reject outliers. This follows from the fact that a relatively large number of optical flow vectors (u, v) will need to be used in determining the values of the parameters, which, in turn, requires a corresponding number of points in both images. In addition, it is difficult to differentiate between optical flow due to rotation around the X and Y axes and translation along the X and Z axes, respectively.

Accordingly, it is desirable to reduce the number of motion parameters to a minimum. The motion of a vehicle 10 along a roadway can be modeled as being constrained to be a translation along the Z axis, as the vehicle 10 moves forward or in reverse, and a rotation around the X and Y axes, as the vehicle 10's path deviates from a straight-line course. In that case, equation (5) reduces to
Image available on "Original document"

If the images are rectified to ensure that the plane of the roadway is parallel to the camera 13's XZ plane, that is, so that the images would be as if the camera 13's optical axis is parallel to the plane of the road, then, in equations (10) a=0 and c=0, in which case
Image available on "Original document"

(11).

In order to rectify the images, the camera 13 will need to be calibrated. A methodology for calibrating the camera 13 and rectifying the images will be described below.

In equations (11) there are three motion parameters, tz (translation along the Z axis), wx (rotation around the X axis) and wy (rotation around the Y axis) to be determined from the flow vectors (u, v) associated with points in at least some portions of the images T and T'. Finding corresponding points in the images T and T', that is, points that are projections of the same point in three-dimensional space in the respective images is based on a"photometric constraint"
I (x,y,t) - I(x + u#t,y+##t,t + #t) = 0 (12),
T which essentially states that the irradiance, or luminance, of the point p = (x, y) in the image T and andthepointp= (x+ uSt, y+ vat) intheimageT',
whichareprojectionsofthesamepoint in three-dimensional space into the respective images, are the same. In practice, equation (12) will not hold exactly because of noise. If, for every point, the noise is modeled as zero mean Gaussian noise, equation 12 reduces to P(I(x,y,t)- I(x+u#t,y+##t,T+#t)) = N (c, 0) (13) and a maximum likelihood can be sought.

Equation (13) can be computationally intensive, and, instead of using that equation, the motion parameters tz, wx and wy can be determined directly from the images by combining the geometric constraints embodied in equation (11) with the photometric constraints embodied in equation (12). In that operation, given two consecutive images

IPR2025-00943
Tesla EX1002 Page 1454

T and T', the goal is to determine the probability
Image available on "Original document"

of that the motion of the vehicle 10 is
Image available on "Original document"

given the two images. The motion m that maximizes (14) is the estimate of the camera 13 motion between the two images T and T', and, thus, the estimate of vehicle 10 motion between the two images.

According to Bayes'rule,
Image available on "Original document"

where
Image available on "Original document"

is the probability that, given image motion m will result in image T', Pt m) is the a priori probability that the motion is m, and
Image available on "Original document"

is the a priori probability that the image is #', It will be assumed that P###, the probability that the motion is #, is uniform in a small region # around the previous motion estimate, that is, the motion estimate generated as between the"i-1 st"and"i-th"images ; it will be appreciated that, if the time period between the times at which images T and T'are recorded is sufficiently small, this assumption will hold. It will further be appreciated that, in equation (15), the denominator P (vu') does not depend on the motion m, and so it does not affect the search for a maximum.

The probability that, given image, motion m will result in image
Image available on "Original document"

be determined by warping the image #' according to the motion #, thereby to generate a warped image #' and determining the sum squared difference ("SSD")
Image available on "Original document"

between corresponding patches, or regions, R in the images that are believed to be a projection of the roadway in the two images. In equation 16,"N"is the number of points in the region R. It will be appreciated that, if the images T and T'are recorded at times"t"and"t+6t, respectively, the warped image 'will represent the image that is assumed would be recorded at time"t"if the motion is m. Using the SSd criteria (equation 16),
Image available on "Original document"

, the probability that image Y" would result from image T and a motion m is given by the probability density function
Image available on "Original document"

where"c"is a normalization factor and"a"is the variance of the noise, which is modeled as a zero mean Gaussian function. Since it is assumed that P###, the probability that the motion is #, is uniform in a small region # around a previous motion estimate, the problem of finding the maximum likelihood motion m for a patch of image T reduces to finding the maximum of the probability density function

IPR2025-00943
Tesla EX1002 Page 1455

Image available on "Original document"

for motion m E M.

Since the motion for which the ego-motion estimation system processor 14 is to generate an estimate is the translational and rotational motion of the vehicle 10 relative to the road, it is desirable for the ego-motion estimation system processor 14 to consider only regions of the images T and T' that comprise projections of the road, and ignore other regions of the images. However, it should be noted that the set R of regions, or patches, of the images that projections of the roadway in the two images # and #', is not known. To accommodate that, instead of attempting to maximize the function defined in equation (18), the image can be tessellated into a set of patches Wj, and a probability density
Image available on "Original document"

generated for each patch using equations (16) and (18) for the respective patches. The probability density over the entire pair of images T and T'will be
Image available on "Original document"

where X, and a, are weighting functions whose values generally reflect the confidence that the"i-th" patch is a projection of the road. The value of the gradient strength Pi for a patch reflects the degree to which the patch the contains a texture, and thus will more likely to contain useful information for use in determining ego motion of the vehicle. The motion m E M for which equation (19) is the maximum will be deemed to correspond to the actual translational and rotational motion of the vehicle 10 as between the locations at which images T and ?'were recorded.

The weighting function Xi for the respective"i-th"patch is generated using patches W, and Wj from respective images T and Y". In determining the values for the weighting functions Xi, for patches W*, W'j that are not of the road, the motion model reflected in equation (11) is not a good fit ; instead, a better fit can be obtained using some other motion of the patch. In addition, for planar image artifacts moving on the roadway surface, such as moving shadows, the maximum of equation (18) will occur far away from the initial guess. Accordingly, the value of the weighting function R for the "i-th" patch Wi, Wi' will correspond to the ratio between the best fit using the motion model in a local region (#) near the initial guess and the best fit using any motion model over large sea2ch region"L."Accordingly, if
Image available on "Original document"

for all # # # is the value for the best fit in a local search region, where Si() denotes the SSD over all points in the "i-th" patch, and
Image available on "Original document"

for all m # L the value for the best fit over all possible image motions, then #i = P1 (22).

P2
Generally, generating the value for P2 (equation 21) can be computationally intensive. To avoid generating P2 according to equation (21), the value for P2 for each patch can be estimated by using the SSD as between a patch in the image # and the correspondingly-positioned patch in the image T', as well as the SSD's as between the patch in the image T and patches translated horizontally and vertically around the correspondingly-positioned patch in the image T', for a selected number of points. That

Tesla EX1002 Page 1456 IPR2025-00943

is, if the patch in image T consists of points p (x, y) centered on p (a, b) (that is, points p (a-a, b-p) through p (a+a, b+p) (a and being integers), with the patch being of dimensions 2a+1 by 2p+I), P2 is generated by using the SSD as between the patch of the same size in image T' consisting of points p (x, y) centered on p (a, b), as well as SSD's as between the patch in image I'and patches of the same size in image T'that are centered on points p (a-6, b-6) through p (a+6, b+6), a total of (26+1) 2 patches in image T'. Each patch in image ?'can be considered as one of the possible image motions. In one embodiment, 8 is selected to be seven, in which case there will be two hundred and twenty five patches in T'for which the SSD will be generated in generating the value for P2.

Unlike weighting function ,, weighting function ai for the respective"i-th"patch is generated using only patch Wi and W ; from image T'. Generally, it should be appreciated that, in threedimensional space, there are three dominant directions for lines, namely, (i) vertical, for buildings, automobiles, and other objects that extend above the roadway surface, (ii) horizontal lines that are generally parallel to the direction of the roadway, and (iii) horizontal lines that are generally perpendicular to the direction of the roadway. In an image T', vertical lines (that is, lines of type (i)) and horizontal lines that are generally perpendicular to the direction of the roadway (that is, lines of type (iii)) will project in to image T'as vertical and horizontal lines, respectively. On the other hand, horizontal lines that are generally parallel to the direction of the roadway (that is, lines of type (ii)) will appear in image T'as lines that pass through a common point, that is, a vanishing point. It will be appreciated that patches W'j in image T'that are projections of the roadway will predominately contain lines of type (ii) and (iii). On the other hand, patches W'that are projections of obstacles, such as automobiles will predominately contain lines of type (i) and (iii), while patches W'that are projections of, for example, buildings, fences, and the like, will contain lines of type (i) and (ii).

Accordingly, the value for weighting function ai for patch Wi will reflect the degree to which it is deemed to contain projections of lines of type (ii) and (iii), and not projections of lines of types (i) and (iii) or types (i) and (ii). Generally, the directions of lines, if any, passing though a patch can be determined in relation to the gradients of the luminance at the various points in the patch Wi.

Each point in the patch W, whose gradient (IX, Iy) is above a selected threshold is considered to lie at or near a line, with the direction of the line being perpendicular to the direction of the gradient.

Thus, for those points, the direction of the line associated therewith can be determined, as can whether the line is of type (i), (ii) or (iii). Thus, for each patch W, in the image T', three sums S" ; (n=i, ii or iii) are generated each corresponding to the sum of the magnitudes of the gradients of the points in the patch that are associated with lines of the respective type, the magnitude corresponding to G= (IXz+IYZ) o. s. A patch W, in image T'is deemed to be : (a) a projection of the roadway if the sums S"i and Siiij are both large and significantly larger than sum Si*, since the sums indicate that the patch is associated with a line or lines that are horizontal and in the direction of the roadway and/or perpendicular thereto, but not a line or lines that are vertical ; (b) be a projection of an obstacle, that is, an object generally in the path of the vehicle, if sums S'i and Siiii are both large and significantly larger than sum Siii since the sums indicate that the patch is associated with a line or lines that are vertical and/or horizontal and perpendicular to the direction of the roadway, but not a line or lines that are horizontal and in the direction of the roadway ; and (c) a projection of an object to the side of the

Tesla EX1002 Page 1457 IPR2025-00943

path of the vehicle if sums S'i and S"i are both large and significantly larger than sum S"',, since the sums indicate that the patch is associated with a line or lines that are vertical and/or horizontal and in the direction of the roadway.
The value of the weighting function ai is assigned to the patch based on the degree to which the patch
W, is deemed to be a projection of the roadway (case (a) above).

It will be appreciated that, if, for a patch W'*, the sum S"i is relatively large, indicating that the patch is associated with a line that is horizontal and in the direction of the roadway, but sums S'i and S..., are relatively small, indicating that the patch is not associated with a line that is vertical or horizontal and perpendicular to the direction of the roadway, it generally cannot be determined from the set of sums generated for the patch whether the patch is a projection of the roadway (case (a) above) or a projection of an object to the side of the path of the vehicle (case (c) above). However, since the patch is not associated with a line that is vertical, it will generally not be deemed to be a projection of an obstacle (case (b) above). In that case, an assessment as to whether the patch is a projection of the roadway (case (a) above) or a projection of an object to the side of the path of the vehicle (case (c) above) by referring to patches adjacent thereto.

As noted above, the value of the gradient strength Pi for a patch reflects the degree to which the patch the contains a texture, and thus will more likely to contain useful information for use in determining ego motion of the vehicle. The gradient strength Pi corresponds to
Image available on "Original document"

For relatively uniform patches, the value of the
Image available on "Original document"

will be relatively low for all motions, in which case the value ofp ; will be relatively low. On the other hand, for patches with texture, the value of the SSD will be relatively high for most motions, in which case the value of p, will be relatively high.

With this background, operations performed by the ego-motion estimation system processor 14 will be describe in connection with the flow chart depicted in FIG. 2. In connection with FIG.

2, it is assumed that the ego-motion estimation system processor 14 already has image T, which it may have used in connection with determining the translational and rotational motion up to the location at which image recorded. With reference to FIG. 2, after the ego-motion estimation system processor 14 has received image T' (step 100), it will rectify the image according to information provided during the camera 13 calibration operation (described below) to provide that the optical axis is parallel to the plane defined by the roadway (step 101). In addition, the egomotion estimation system processor 14 will generate an initial guess as to the translational and rotational motion, using the previous motion estimate and, perhaps information from other sensors if available (step 102). For example, the ego-motion estimation system processor 14 may make use of information from the vehicle 10's speedometer, as well as information as to the time period between the time at which image T was recorded and the time at which image T'was recorded, in generating the initial guess. Generally, it will be appreciated that the time period will be fixed, and will preferably the same for each

IPR2025-00943
Tesla EX1002 Page 1458

successive pair of images T and T'. After the ego-motion estimation system processor 14 has generated the initial guess, it will use the initial guess to warp image T'toward image T, thereby to generate a warped image (step 103).

After the ego-motion estimation system processor 14 has generated the warped image (step 103), it will divide the image # and the warped image #' into patches and, for each pair of corresponding patches in the two images T and', generate the weighting value Xi. In that operation, the ego-motion estimation system processor 14 will select a patch in the image T (step 104) and generate values for P2 (step 105), P, (equation 20) (step 106) and Xj (equation 22) (step 107) as described above. In addition, the ego-motion estimation system processor 14 can generate the value for Pi (equation 23) and ai (step 108). After the ego-motion estimation system processor 14 has generated performed steps 105 through 108 for the selected patch, it will determine whether all of the patches in image T have been processed (step 109) and if not, return to step 104 to select another patch and perform steps 105 through 109 in connection therewith.

The ego-motion estimation system processor 14 will perform steps 104 through 109 in connection with each patch in the image T.. After the ego-motion estimation system processor 14 has performed steps 104 through 109 in connection with all of the patches in the image, it will sequence from step 109 to step 110 to search for the motion m that maximizes the value provided by equation (19) (step 110). That motion m will comprise values for translation tz and rotation wx, wy parameters that will constitute the estimate of the motion of the vehicle 10 as between the point in time at which image T was recorded and the point in time at which image T'is recorded. The ego-motion estimation system processor 14 can perform operations described above in connection with each successive pair of images T and T'to estimate the motion of the vehicle 10.

In performing steps 106 (to generate the values for P_i) and 110 (to determine the motion m that maximizes the value provided by equation (19)), the ego-motion estimation system processor
14 can perform a gradient descent that is limited to a selected cube-shaped region around the initial guess.

In determining the initial guess (step 102) for each new image T', the ego-motion estimation system processor 14 can use the estimate of the motion generated for the previously-received image.

In addition, the size of the region # can be adjusted adaptively.

A second methodology for determining the ego-motion of a vehicle 10 will be described in connection with FIG. 3. As with the methodology described above in connection with FIG. 2, it is assumed that the images have been rectified so that the image planes are perpendicular to the plane represented by the roadway on which the vehicle 10 is traveling, and their horizontal ("x") axes are parallel to the plane represented by the roadway. By way of background, in that case, the equation of a plane AX+ BY+ CZ = 1 (24) (reference equation (3)) reduces to
BY = 1 (25), in which case
image available on "Original document"

is the height of the optical axis of the camera 13 (or, more specifically, the Z axis) above the road. Since, for a point p (x, y) in an image that is a projection of a

Tesla EX1002 Page 1459 IPR2025-00943

point P (X, Y, Z) y in three-dimensional space, y = f-, equation (25) becomes
1 -= boy (26)
The brightness constraint is ulX+vly+l=0 (27) for each point, where, at each point (x, y) in the image, Ix and Iy are the horizontal and vertical components of the spatial gradient of the luminance and It is the time gradient of the luminance. In addition, the equations for the components (u, v) of the flow vector (reference equation (1)) can be written
Image available on "Original document"

where [w] x is a skew-symmetric matrix, in which case
Image available on "Original document"

where
Image available on "Original document"

For motion constrained to a plane, equation (29) reduces to (ax + by + c) STt + ST[w} xp + It = 0
Since the images are rectified, equation (26) will hold. In addition, since rotational motion is constrained so as to occur only around the vertical ("Y") axis, wx=wz=0. Accordingly, equation (31) will correspond to
Image available on "Original document"

Expanding equation 32, and using equation (30),
Image available on "Original document"

Generally, the portion uwy of the vertical component of the flow vector that arises due to rotation (wy) of the vehicle 10 will be larger than the portion ux that arises due to translation (tx) along the "X"axis, and the portions will have a very different form. The portion ux that arises due to translation along the"X"axis will depend on the distance that the vehicle 10 moves between the times the images are recorded, and will be larger for points towards the bottom of the image than for points further up. In that case, expressions for uwy and ux are
Image available on "Original document"

Since, for small rotations around the vertical"Y,"axis, the angle of rotation wy is approximately t wy ~ x (36), tz
Image available on "Original document"

It will be appreciated that values for ty, the component of the translation t in the vertical direction, and wx and wz, the X and Z components of rotation w, will be zero. Accordingly, after the egomotion estimation system processor 14 receives a new image T', it will determine the values for t, and t3, the components of the translation t in the forward (along the Z axis) and side (along the X axis) directions, and wy, the component of rotation around the vertical (Y) axis. In that operation, the ego-motion estimation system processor 14 will generate an initial estimate as to the motion (step 150, FIG. 3) and use that initial estimate to generate a warped image sus (step 151). The egomotion estimation system processor 14 can use information from a number of

IPR2025-00943
Tesla EX1002 Page 1460

sources in connection with generating the initial estimate (step 150), including information from, for example, the vehicle 10's speedometer. Thereafter, the ego-motion estimation system processor 14 divides the image T and Y'into patches (step 152) and determines which patches are likely to comprise images of the roadway (step 153). In determining which patches are likely to comprise images of the roadway (reference step 153), the ego-motion estimation system processor 14 can generate an SSD (equation 16) for corresponding patches in the images T and V"and deten-nine that patches in the warped image F'that comprise images of the roadway will be those patches with a relatively high SSD value. The ego-motion estimation system processor 14 then uses the patches identified in step 153 to minimize a cost function of the form image available on "Original document"

where W, and Wz are weighting matrices that essentially describe confidence in the priors or smoothness values (step 154). Equation (38) can be formalized in the form of a Kalman filter, and the value of'p"can be selected to be one or two depending on whether the L, or L2 norm is to be used.

As noted above, for methodologies described above in connection with both FIGS. 2 and 3, the ego-motion estimation system processor 14 will initially rectify the images as received from the camera 13. In the above description, the images I and I'are images as rectified by the ego-motion estimation syst%n processor 14. To rectify the image3 as received from the camera 13, the camera 13 will need to be calibrated during a calibration operation prior to use in connection with recording images for use in estimating vehicle 10 motion as described above. Before describing operations to be performed during calibration, it would be helpful to consider the effects of incorrect calibration.

If, for example, a vehicle is driving along a straight road, and if the optical axis of the camera is aligned with the direction of motion, the flow field of successively recorded images will be an expansion field with the focus of expansion located at the center of the respective images, that is, at the origin (x, y) = (0, 0) of the image's rectilinear coordinate system. On the other hand, if the camera is mounted on the vehicle with a small rotation around the vertical ("Y") axis in three-dimensional space, then the focus of expansion will be displaced along the image's horizontal ("x") axis. In that case, the motion model defined by equation (11) will not account for the flow field, but will be well approximated by a forward translation and a rotational velocity wy around the vertical ("Y") axis.

Accordingly, errors in the orientation of the camera around the vertical ("Y") axis in three dimensional space will create a bias in the rotation estimate, in which case a curved path would be estimated instead of a straight path. Similarly, errors in the camera's orientation around the horizontal ("X") axis in three dimensional space will cause a bias in the pitch estimate. Based on these observations, a calibration operation can be performed by having the camera record a sequence of images while the vehicle is being driven down a straight roadway. The vehicle's ego-motion is estimated as described above in connection with FIGS. 2 or 3, and calibration parameters are estimated that would cause the ego-motion to integrate into a straight path.

The invention provides a number of advantages. In particular, the invention provides an arrangement for determining ego-motion of a vehicle 10 on a roadway from a series of images recorded by a camera 13 mounted on the vehicle 10, at least a portion of the images comprising projections of the roadway, and without requiring mechanical

Tesla EX1002 Page 1461 IPR2025-00943

sensors which are normally not provided with a vehicle 10 and that would, if provided, increase the cost and maintenance expenses thereof.

It will be appreciated that a system in accordance with the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof, any portion of which may be controlled by a suitable program. Any program may in whole or in part comprise part of or be stored on the system in a conventional manner, or it may in whole or in part be provided in to the system over a network or other mechanism for transferring information in a conventional manner. In addition, it will be appreciated that the system may be operated and/or otherwise controlled by means of information provided by an operator using operator input elements (not shown) which may be connected directly to the system or which may transfer the information to the system over a network or other mechanism for transferring information in a conventional manner.

The foregoing description has been limited to a specific embodiment ofthis invention. It will be apparent, however, that various variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention. It is the object of the appended claims to cover these and such other variations and modifications as come within the true spirit and scope of the invention.

What is claimed as new and desired to be secured by Letters Patent of the United States is :

Espacenet

Claims: EP1257971 (A2) — 2002-11-20

SYSTEM AND METHOD FOR ESTIMATING EGO-MOTION OF A MOVING VEHICLE USING SUCCESSIVE IMAGES RECORDED ALONG THE VEHICLE'S PATH OF MOTION

Claims not available for EP1257971 (A2)
Claims of corresponding document: WO0139120 (A2)

A high quality text as facsimile in your desired language may be available amongst the following family members:

CA2392652 (A1) WO0139120 (A2)

- Original claims
- Claims tree

The EPO does not accept any responsibility for the accuracy of data and information originating from other authorities than the EPO; in particular, the EPO does not guarantee that they are complete, up-to-date or fit for specific purposes.

CLAIMS 1. An ego-motion determination system for generating an estimate as to the ego-motion of a vehicle moving along a roadway comprising :
A. an image information receiver configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway ; and
B. a processor configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Veröffentlichungsnummer:
(11) Publication number: **EP 1 257 971 A0**
(11) Numéro de publication:

Internationale Anmeldung veröffentlicht durch die Weltorganisation für geistiges Eigentum unter der Nummer:

WO 01/039120 (art. 158 des EPÜ).

International application published by the World Intellectual Property Organisation under number:

WO 01/039120 (art. 158 of the EPC).

Demande internationale publiée par l'Organisation Mondiale de la Propriété sous le numéro:

WO 01/039120 (art. 158 de la CBE).

Espacenet

Bibliographic data: EP1257971 (A4) — 2005-07-06

SYSTEM AND METHOD FOR ESTIMATING EGO-MOTION OF A MOVING VEHICLE USING SUCCESSIVE IMAGES RECORDED ALONG THE VEHICLE'S PATH OF MOTION

Inventor(s): STEIN GIDEON [IL]; SHASHUA AMNON [IL]; MANO OFER [IL] ± (STEIN, GIDEON, ; SHASHUA, AMNON, ; MANO, OFER)

Applicant(s): MOBILEYE INC [US] ± (MOBILEYE, INC)

Classification: **- international:** ***G01S11/12; G05D1/02; G06T1/00; G06T7/20; G08G1/16;*** *G01S5/16;* (IPC1-7): G05D1/02; G06T7/20

- cooperative: **G01S11/12 (EP);** G01S5/163 (EP)

Application number: EP20000980706 20001127 Global Dossier

Priority number (s): WO2000US32143 20001127 ; US19990167587P 19991126 ; US20000230166P 20000901

Also published as: EP1257971 (A2) AU1793301 (A) CA2392652 (A1) JP2003515827 (A) WO0139120 (A2) more

Abstract not available for EP1257971 (A4)
Abstract of corresponding document: WO0139120 (A2)



An ego-motion determination system is disclosed for generating an estimate as to the ego-motion of a vehicle moving along a roadway. The ego-motion determination system includes an image information receiver and a processor. The image information receiver is configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway. The processor is configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle, including the translation of the vehicle in the forward direction and the rotation of the vehicle around a vertical axis as between, for example, successive images.;An ego-motion determination system (fig. 1, item 12) is disclosed for generating an estimate as to the ego-motion of a vehicle (fig. 1, item 10) moving along a roadway (fig. 1, item 11). The ego-motion determination system includes an image information receiver (fig. 1, item 13) and a processor (fig. 1, item 14). The image information receiver (fig. 1, item 13) is configured to receive image information relating

to a series of at least two images recorded as the vehicle moves along a roadway (fig. 1, item 11). The processor is configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle, including the translation of the vehicle in the forward direction and the rotation of the vehicle around a vertical axis as between, for example, successive images (see fig. 3)

IPR2025-00943
Tesla EX1002 Page 1466

Espacenet

Description: EP1257971 (A4) — 2005-07-06

SYSTEM AND METHOD FOR ESTIMATING EGO-MOTION OF A MOVING VEHICLE USING SUCCESSIVE IMAGES RECORDED ALONG THE VEHICLE'S PATH OF MOTION

Description not available for EP1257971 (A4)
Description of corresponding document: WO0139120 (A2)

A high quality text as facsimile in your desired language may be available amongst the following family members:

CA2392652 (A1) WO0139120 (A2)

The EPO does not accept any responsibility for the accuracy of data and information originating from other authorities than the EPO; in particular, the EPO does not guarantee that they are complete, up-to-date or fit for specific purposes.

SYSTEM AND METHOD FOR ESTIMATING EGO-MOTION OF A MOVING VEHICLE USING
SUCCESSIVE IMAGES RECORDE ALONG THE VEHICLE'S PATH OF MOTION
FIELD OF THE INVENTION
The invention relates generally to the field of systems and methods for estimating egomotion (that is,"self-motion) of a moving vehicle, and more specifically to systems and methods that estimate ego-motion using successively-recorded images recorded along the vehicle's path of motion.

BACKGROUND OF THE INVENTION
Accurate estimation of the ego- ("self-") motion of a vehicle relative to a roadway is an important component in autonomous driving and computer vision-based driving assistance. Using computer vision techniques to provide assistance while driving, instead of mechanical sensors, allows for the use of the information that is recorded for use in estimating vehicle movement to also be used in detecting obstacles, identifying lanes and the like, without the need for calibration between sensors as would be necessary with mechanical sensors. This reduces cost and maintenance.

There are several problems in estimating ego-motion of a vehicle. Typically, roads have few feature points, if any. The most obvious features in a road, such as lane markings, have a generally linear structure, whereas background image structures, such as those associated with other vehicles, buildings, trees, and the like, will typically have many feature points. This will make image-or optical-flow-based estimation difficult in

practice. In addition, typically images that are recorded for ego-motion estimation will contain a large amount of"outlier"information that is either not useful in estimating ego-motion, or that may result in poor estimation. For example, in estimating of ego-motion relative to a fixed structure such as a road, images of objects that are moving relative to the road, such as other moving vehicles and even moving wipers, will contribute false information for the ego-motion estimation. In addition, conditions that degrade image quality, such as raindrops and glare, will also make accurate ego-motion estimation difficult.

SUMMARY OF THE INVENTION

The invention provides a new and improved system and method for estimating ego-motion using successively-recorded images recorded along the vehicle's path of motion..

In brief summary, the invention provides an ego-motion determination system for generating an estimate as to the ego-motion of a vehicle moving along a roadway. The ego-motion determination system includes an image information receiver and a processor. The image information receiver is configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway. The processor is configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle, including the translation of the vehicle in the forward direction and the rotation of the vehicle around a vertical axis as between, for example, successive images.

Several methodologies are disclosed for generating the ego-motion estimate of the vehicle.

BRIEF DESCRIPTION OF THE DRAWINGS

This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which :
FIG. 1 schematically depicts a vehicle moving on a roadway and including an ego-motion estimation system constructed in accordance with the invention ;
FIG. 2 is a flow chart depicting operations performed by the ego-motion estimation system in determining ego-motion of the vehicle in accordance with one methodology ; and
FIG. 3 is a flow chart depicting operations performed by the ego-motion estimation system in determining ego-motion of the vehicle in accordance with a second methodology.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

FIG. 1 schematically depicts a vehicle 10 moving on a roadway 11 and including an egomotion estimation system 12 constructed in accordance with the invention. The vehicle 10 may be any kind of vehicle 10 that may move on the roadway 11, including, but not limited to automobiles, trucks, buses and the like. The ego-motion estimation system 12 includes a camera 13 and a egomotion estimation system processor 14. The camera 13 is mounted on the vehicle 10 and is preferably pointed in a forward direction, that is, in the direction in which the vehicle would normally move, to record successive images as the vehicle moves over the roadway. Preferably as the camera 13 records each image, it will provide the image to the ego-motion estimation system processor 14. The ego-motion estimation system processor 14, in turn, will process information that it obtains from the successive images, possibly along with other

IPR2025-00943
Tesla EX1002 Page 1468

information, such as information from the vehicle's speedometer (not separately shown) to determine the ego-motion (that is, the selfmotion) of the vehicle relative to the roadway 11. The ego-motion estimation system processor 14 may also be mounted in or on the vehicle 11 and may form part thereof. The ego-motion estimates generated by the ego-motion estimation system processor 14 may be used for a number of things, including, but not limited to obstacle and lane detection, autonomous driving by the vehicle, perhaps also using positioning information from, for example, the global positioning system ("GPS") and roadway mapping information from a number of sources known to those skilled in the art, and the like. Operations performed by the ego-motion estimation system processor 14 in determining egomotion of the vehicle 10 will be described in connection with the flow charts depicted in FIGS. 2 and 3.

Before proceeding further, it would be helpful to provide background of the operations performed by the ego-motion estimation system processor 14 depicted in FIG. 1. Generally, as between two images T and T'the ego-motion estimation system processor 14 attempts to determine the translation t = (tx, ty, tz) and rotation w = (WX, wy, wz) (where"T"refers to the transpose operation, t, refers to translation along the respective"X,""Y"and"Z"axes, and w, refers to rotation around the respective axis) of the camera 13 affixed to the vehicle 10. Since the camera 13 is affixed to the vehicle 10, the translation and rotation of the camera 13 will also conform to the translation and rotation of the vehicle 10. In that case, for a point p = (x,y)T (where "x" and "y" are coordinates of apoint orfeature inthe image) thatis aprojection of apoint P = (X, Y, Z) (where "X,""Y"and"Z"are coordinates of the point in three-dimensional space), the flow vector for the point, that is, the vector indicating the motion of the same point in three dimensional space from is position in the image T to the image S", has components (u, v)
Image available on "Original document"

where"X"in equations (1) refers to the matrix cross-product and
Image available on "Original document"

where"f"is the focal length of the camera 13, which is presumed to be known. The roadway on which the vehicle 10 is traveling is modeled as a plane. The equation for points on a plane is AX+ BY+ BZ = 1 (3), where the"X"and"Y"axes correspond to horizontal"x"and vertical"y"axes of the image plane, and the"Z"axis is perpendicular to the image plane. The camera 13 may be tilted slightly downwardly to increase the amount of the image that corresponds to the road, and reduce the amount of the image that corresponds to other features, such as other traffic, buildings, trees and the like, and in that case, the images T and T'will be rectified so that the images'z axis will be parallel to the plane of the roadway, as will be described below.

Dividing equation (3) by "Z" provides
1/Z = ax + by + c (4), where a=Alf, b=Blf and c=C. Substituting equation (4) into equations (1) results in
Image available on "Original document"

Expanding equations (5) results in
Image available on "Original document"

Equations (6) and (7) are a special case (the"calibrated camera 13"case) of an eight

parameter model for a camera 13 moving relative to a plane : U = alx+ a2y+ a3+ a7x 2 + a8xy (8) v-a4x+ asy+ a6 + axy+ agyz (9)
Given the flow vector (u, v), one can recover the parameters a,, i=1,..., 8, from which one can recover the motion parameters t and w.

One problem arises in connection with the methodology described above in connection with equations (1) through (9), namely, given the large number of parameters ai whose values need to be determined, it is difficult to devise a satisfactory method that will reject outliers. This follows from the fact that a relatively large number of optical flow vectors (u, v) will need to be used in determining the values of the parameters, which, in turn, requires a corresponding number of points in both images. In addition, it is difficult to differentiate between optical flow due to rotation around the X and Y axes and translation along the X and Z axes, respectively.

Accordingly, it is desirable to reduce the number of motion parameters to a minimum. The motion of a vehicle 10 along a roadway can be modeled as being constrained to be a translation along the Z axis, as the vehicle 10 moves forward or in reverse, and a rotation around the X and Y axes, as the vehicle 10's path deviates from a straight-line course. In that case, equation (5) reduces to
Image available on "Original document"

If the images are rectified to ensure that the plane of the roadway is parallel to the camera 13's XZ plane, that is, so that the images would be as if the camera 13's optical axis is parallel to the plane of the road, then, in equations (10) a=0 and c=0, in which case
Image available on "Original document"

(11).

In order to rectify the images, the camera 13 will need to be calibrated. A methodology for calibrating the camera 13 and rectifying the images will be described below.

In equations (11) there are three motion parameters, tz (translation along the Z axis), wx (rotation around the X axis) and wy (rotation around the Y axis) to be determined from the flow vectors (u, v) associated with points in at least some portions of the images T and T'. Finding corresponding points in the images T and T', that is, points that are projections of the same point in three-dimensional space in the respective images is based on a"photometric constraint"
I (x,y,t) - I(x + u#t,y+##t,t + #t) = 0 (12),
T which essentially states that the irradiance, or luminance, of the point p = (x, y) in the image T and andthepointp= (x+ uSt, y+ vat) intheimageT',
whichareprojectionsofthesamepoint in three-dimensional space into the respective images, are the same. In practice, equation (12) will not hold exactly because of noise. If, for every point, the noise is modeled as zero mean Gaussian noise, equation 12 reduces to P(I(x,y,t)- I(x+u#t,y+##t,T+#t)) = N (c, 0) (13) and a maximum likelihood can be sought.

Equation (13) can be computationally intensive, and, instead of using that equation, the motion parameters tz, wx and wy can be determined directly from the images by combining the geometric constraints embodied in equation (11) with the photometric constraints embodied in equation (12). In that operation, given two consecutive images

IPR2025-00943
Tesla EX1002 Page 1470

T and T', the goal is to determine the probability
Image available on "Original document"

of that the motion of the vehicle 10 is
Image available on "Original document"

given the two images. The motion m that maximizes (14) is the estimate of the camera 13 motion between the two images T and T', and, thus, the estimate of vehicle 10 motion between the two images.

According to Bayes'rule,
Image available on "Original document"

where
Image available on "Original document"

is the probability that, given image motion m will result in image T', Pt m) is the a priori probability that the motion is m, and
Image available on "Original document"

is the a priori probability that the image is #', It will be assumed that P###, the probability that the motion is #, is uniform in a small region # around the previous motion estimate, that is, the motion estimate generated as between the"i-1 st"and"i-th"images ; it will be appreciated that, if the time period between the times at which images T and T'are recorded is sufficiently small, this assumption will hold. It will further be appreciated that, in equation (15), the denominator P (vu') does not depend on the motion m, and so it does not affect the search for a maximum.

The probability that, given image, motion m will result in image
Image available on "Original document"

be determined by warping the image #' according to the motion #, thereby to generate a warped image #' and determining the sum squared difference ("SSD")
Image available on "Original document"

between corresponding patches, or regions, R in the images that are believed to be a projection of the roadway in the two images. In equation 16,"N"is the number of points in the region R. It will be appreciated that, if the images T and T'are recorded at times"t"and"t+6t, respectively, the warped image 'will represent the image that is assumed would be recorded at time"t"if the motion is m. Using the SSd criteria (equation 16),
Image available on "Original document"

, the probability that image Y" would result from image T and a motion m is given by the probability density function
Image available on "Original document"

where"c"is a normalization factor and"a"is the variance of the noise, which is modeled as a zero mean Gaussian function. Since it is assumed that P###, the probability that the motion is #, is uniform in a small region # around a previous motion estimate, the problem of finding the maximum likelihood motion m for a patch of image T reduces to finding the maximum of the probability density function

IPR2025-00943
Tesla EX1002 Page 1471

Image available on "Original document"

for motion m E M.

Since the motion for which the ego-motion estimation system processor 14 is to generate an estimate is the translational and rotational motion of the vehicle 10 relative to the road, it is desirable for the ego-motion estimation system processor 14 to consider only regions of the images T and T' that comprise projections of the road, and ignore other regions of the images. However, it should be noted that the set R of regions, or patches, of the images that projections of the roadway in the two images # and #', is not known. To accommodate that, instead of attempting to maximize the function defined in equation (18), the image can be tessellated into a set of patches Wj, and a probability density
Image available on "Original document"

generated for each patch using equations (16) and (18) for the respective patches. The probability density over the entire pair of images T and T'will be
Image available on "Original document"

where X, and a, are weighting functions whose values generally reflect the confidence that the"i-th" patch is a projection of the road. The value of the gradient strength Pi for a patch reflects the degree to which the patch the contains a texture, and thus will more likely to contain useful information for use in determining ego motion of the vehicle. The motion m E M for which equation (19) is the maximum will be deemed to correspond to the actual translational and rotational motion of the vehicle 10 as between the locations at which images T and ?'were recorded.

The weighting function Xi for the respective"i-th"patch is generated using patches W, and Wj from respective images T and Y". In determining the values for the weighting functions Xi, for patches W*, W'j that are not of the road, the motion model reflected in equation (11) is not a good fit ; instead, a better fit can be obtained using some other motion of the patch. In addition, for planar image artifacts moving on the roadway surface, such as moving shadows, the maximum of equation (18) will occur far away from the initial guess. Accordingly, the value of the weighting function R for the "i-th" patch Wi, Wi' will correspond to the ratio between the best fit using the motion model in a local region (#) near the initial guess and the best fit using any motion model over large search region"L."Accordingly, if
Image available on "Original document"

for all # # # is the value for the best fit in a local search region, where Si() denotes the SSD over all points in the "i-th" patch, and
Image available on "Original document"

for all m # L the value for the best fit over all possible image motions, then #i = P1 (22).

P2
Generally, generating the value for P2 (equation 21) can be computationally intensive. To avoid generating P2 according to equation (21), the value for P2 for each patch can be estimated by using the SSD as between a patch in the image # and the correspondingly-positioned patch in the image T', as well as the SSD's as between the patch in the image T and patches translated horizontally and vertically around the correspondingly-positioned patch in the image T', for a selected number of points. That

Tesla EX1002 Page 1472 IPR2025-00943

is, if the patch in image T consists of points p (x, y) centered on p (a, b) (that is, points p (a-a, b-p) through p (a+a, b+p) (a and being integers), with the patch being of dimensions 2a+1 by 2p+I), P2 is generated by using the SSD as between the patch of the same size in image T' consisting of points p (x, y) centered on p (a, b), as well as SSD's as between the patch in image I'and patches of the same size in image T'that are centered on points p (a-6, b-6) through p (a+6, b+6), a total of (26+1) 2 patches in image T'. Each patch in image ?'can be considered as one of the possible image motions. In one embodiment, 8 is selected to be seven, in which case there will be two hundred and twenty five patches in T'for which the SSD will be generated in generating the value for P2.

Unlike weighting function ,, weighting function ai for the respective"i-th"patch is generated using only patch Wi and W ; from image T'. Generally, it should be appreciated that, in threedimensional space, there are three dominant directions for lines, namely, (i) vertical, for buildings, automobiles, and other objects that extend above the roadway surface, (ii) horizontal lines that are generally parallel to the direction of the roadway, and (iii) horizontal lines that are generally perpendicular to the direction of the roadway. In an image T', vertical lines (that is, lines of type (i)) and horizontal lines that are generally perpendicular to the direction of the roadway (that is, lines of type (iii)) will project in to image T'as vertical and horizontal lines, respectively. On the other hand, horizontal lines that are generally parallel to the direction of the roadway (that is, lines of type (ii)) will appear in image T'as lines that pass through a common point, that is, a vanishing point. It will be appreciated that patches W'j in image T'that are projections of the roadway will predominately contain lines of type (ii) and (iii). On the other hand, patches W'that are projections of obstacles, such as automobiles will predominately contain lines of type (i) and (iii), while patches W'that are projections of, for example, buildings, fences, and the like, will contain lines of type (i) and (ii).

Accordingly, the value for weighting function ai for patch Wi will reflect the degree to which it is deemed to contain projections of lines of type (ii) and (iii), and not projections of lines of types (i) and (iii) or types (i) and (ii). Generally, the directions of lines, if any, passing though a patch can be determined in relation to the gradients of the luminance at the various points in the patch Wi.

Each point in the patch W, whose gradient (IX, Iy) is above a selected threshold is considered to lie at or near a line, with the direction of the line being perpendicular to the direction of the gradient.

Thus, for those points, the direction of the line associated therewith can be determined, as can whether the line is of type (i), (ii) or (iii). Thus, for each patch W, in the image T', three sums S" ; (n=i, ii or iii) are generated each corresponding to the sum of the magnitudes of the gradients of the points in the patch that are associated with lines of the respective type, the magnitude corresponding to G= (IXz+IYZ) o. s. A patch W, in image T'is deemed to be : (a) a projection of the roadway if the sums S"i and Siiij are both large and significantly larger than sum Si*, since the sums indicate that the patch is associated with a line or lines that are horizontal and in the direction of the roadway and/or perpendicular thereto, but not a line or lines that are vertical ; (b) be a projection of an obstacle, that is, an object generally in the path of the vehicle, if sums S'i and Siiii are both large and significantly larger than sum Siii since the sums indicate that the patch is associated with a line or lines that are vertical and/or horizontal and perpendicular to the direction of the roadway, but not a line or lines that are horizontal and in the direction of the roadway ; and (c) a projection of an object to the side of the

Tesla EX1002 Page 1473 IPR2025-00943

path of the vehicle if sums S'i and S"i are both large and significantly larger than sum S"',, since the sums indicate that the patch is associated with a line or lines that are vertical and/or horizontal and in the direction of the roadway.
The value of the weighting function ai is assigned to the patch based on the degree to which the patch
W, is deemed to be a projection of the roadway (case (a) above).

It will be appreciated that, if, for a patch W'*, the sum S"i is relatively large, indicating that the patch is associated with a line that is horizontal and in the direction of the roadway, but sums S'i and S..., are relatively small, indicating that the patch is not associated with a line that is vertical or horizontal and perpendicular to the direction of the roadway, it generally cannot be determined from the set of sums generated for the patch whether the patch is a projection of the roadway (case (a) above) or a projection of an object to the side of the path of the vehicle (case (c) above). However, since the patch is not associated with a line that is vertical, it will generally not be deemed to be a projection of an obstacle (case (b) above). In that case, an assessment as to whether the patch is a projection of the roadway (case (a) above) or a projection of an object to the side of the path of the vehicle (case (c) above) by referring to patches adjacent thereto.

As noted above, the value of the gradient strength Pi for a patch reflects the degree to which the patch the contains a texture, and thus will more likely to contain useful information for use in determining ego motion of the vehicle. The gradient strength Pi corresponds to
Image available on "Original document"

For relatively uniform patches, the value of the
Image available on "Original document"

will be relatively low for all motions, in which case the value ofp ; will be relatively low. On the other hand, for patches with texture, the value of the SSD will be relatively high for most motions, in which case the value of p, will be relatively high.

With this background, operations performed by the ego-motion estimation system processor 14 will be describe in connection with the flow chart depicted in FIG. 2. In connection with FIG.

2, it is assumed that the ego-motion estimation system processor 14 already has image T, which it may have used in connection with determining the translational and rotational motion up to the location at which image recorded. With reference to FIG. 2, after the ego-motion estimation system processor 14 has received image T' (step 100), it will rectify the image according to information provided during the camera 13 calibration operation (described below) to provide that the optical axis is parallel to the plane defined by the roadway (step 101). In addition, the egomotion estimation system processor 14 will generate an initial guess as to the translational and rotational motion, using the previous motion estimate and, perhaps information from other sensors if available (step 102). For example, the ego-motion estimation system processor 14 may make use of information from the vehicle 10's speedometer, as well as information as to the time period between the time at which image T was recorded and the time at which image T'was recorded, in generating the initial guess. Generally, it will be appreciated that the time period will be fixed, and will preferably the same for each

IPR2025-00943
Tesla EX1002 Page 1474

successive pair of images T and T'. After the ego-motion estimation system processor 14 has generated the initial guess, it will use the initial guess to warp image T'toward image T, thereby to generate a warped image (step 103).

After the ego-motion estimation system processor 14 has generated the warped image (step 103), it will divide the image # and the warped image #' into patches and, for each pair of corresponding patches in the two images T and', generate the weighting value Xi. In that operation, the ego-motion estimation system processor 14 will select a patch in the image T (step 104) and generate values for P2 (step 105), P, (equation 20) (step 106) and Xj (equation 22) (step 107) as described above. In addition, the ego-motion estimation system processor 14 can generate the value for Pi (equation 23) and ai (step 108). After the ego-motion estimation system processor 14 has generated performed steps 105 through 108 for the selected patch, it will determine whether all of the patches in image T have been processed (step 109) and if not, return to step 104 to select another patch and perform steps 105 through 109 in connection therewith.

The ego-motion estimation system processor 14 will perform steps 104 through 109 in connection with each patch in the image T.. After the ego-motion estimation system processor 14 has performed steps 104 through 109 in connection with all of the patches in the image, it will sequence from step 109 to step 110 to search for the motion m that maximizes the value provided by equation (19) (step 110). That motion m will comprise values for translation tz and rotation wx, wy parameters that will constitute the estimate of the motion of the vehicle 10 as between the point in time at which image T was recorded and the point in time at which image T'is recorded. The ego-motion estimation system processor 14 can perform operations described above in connection with each successive pair of images T and T'to estimate the motion of the vehicle 10.

In performing steps 106 (to generate the values for P_i) and 110 (to determine the motion m that maximizes the value provided by equation (19)), the ego-motion estimation system processor
14 can perform a gradient descent that is limited to a selected cube-shaped region around the initial guess.

In determining the initial guess (step 102) for each new image T', the ego-motion estimation system processor 14 can use the estimate of the motion generated for the previously-received image.

In addition, the size of the region # can be adjusted adaptively.

A second methodology for determining the ego-motion of a vehicle 10 will be described in connection with FIG. 3. As with the methodology described above in connection with FIG. 2, it is assumed that the images have been rectified so that the image planes are perpendicular to the plane represented by the roadway on which the vehicle 10 is traveling, and their horizontal ("x") axes are parallel to the plane represented by the roadway. By way of background, in that case, the equation of a plane AX+ BY+ CZ = 1 (24) (reference equation (3)) reduces to
BY = 1 (25), in which case
image available on "Original document"

is the height of the optical axis of the camera 13 (or, more specifically, the Z axis) above the road. Since, for a point p (x, y) in an image that is a projection of a

Tesla EX1002 Page 1475 IPR2025-00943

point P (X, Y, Z) y in three-dimensional space, y = f-, equation (25) becomes
1 -= boy (26)
The brightness constraint is ulX+vly+l=0 (27) for each point, where, at each point (x, y) in the image, Ix and Iy are the horizontal and vertical components of the spatial gradient of the luminance and It is the time gradient of the luminance. In addition, the equations for the components (u, v) of the flow vector (reference equation (1)) can be written
Image available on "Original document"

where [w] x is a skew-symmetric matrix, in which case
Image available on "Original document"

where
Image available on "Original document"

For motion constrained to a plane, equation (29) reduces to (ax + by + c) STt + ST[w} xp + It = 0
Since the images are rectified, equation (26) will hold. In addition, since rotational motion is constrained so as to occur only around the vertical ("Y") axis, wx=wz=0. Accordingly, equation (31) will correspond to
Image available on "Original document"

Expanding equation 32, and using equation (30),
Image available on "Original document"

Generally, the portion uwy of the vertical component of the flow vector that arises due to rotation (wy) of the vehicle 10 will be larger than the portion ux that arises due to translation (tx) along the "X"axis, and the portions will have a very different form. The portion ux that arises due to translation along the"X"axis will depend on the distance that the vehicle 10 moves between the times the images are recorded, and will be larger for points towards the bottom of the image than for points further up. In that case, expressions for uwy and ux are
Image available on "Original document"

Since, for small rotations around the vertical"Y,"axis, the angle of rotation wy is approximately t wy ~ x (36), tz
Image available on "Original document"

It will be appreciated that values for ty, the component of the translation t in the vertical direction, and wx and wz, the X and Z components of rotation w, will be zero. Accordingly, after the egomotion estimation system processor 14 receives a new image T', it will determine the values for t, and t3, the components of the translation t in the forward (along the Z axis) and side (along the X axis) directions, and wy, the component of rotation around the vertical (Y) axis. In that operation, the ego-motion estimation system processor 14 will generate an initial estimate as to the motion (step 150, FIG. 3) and use that initial estimate to generate a warped image sus (step 151). The egomotion estimation system processor 14 can use information from a number of

IPR2025-00943
Tesla EX1002 Page 1476

sources in connection with generating the initial estimate (step 150), including information from, for example, the vehicle 10's speedometer. Thereafter, the ego-motion estimation system processor 14 divides the image T and Y'into patches (step 152) and determines which patches are likely to comprise images of the roadway (step 153). In determining which patches are likely to comprise images of the roadway (reference step 153), the ego-motion estimation system processor 14 can generate an SSD (equation 16) for corresponding patches in the images T and V"and deten-nine that patches in the warped image F'that comprise images of the roadway will be those patches with a relatively high SSD value. The ego-motion estimation system processor 14 then uses the patches identified in step 153 to minimize a cost function of the form image available on "Original document"

where W, and Wz are weighting matrices that essentially describe confidence in the priors or smoothness values (step 154). Equation (38) can be formalized in the form of a Kalman filter, and the value of'p"can be selected to be one or two depending on whether the L, or L2 norm is to be used.

As noted above, for methodologies described above in connection with both FIGS. 2 and 3, the ego-motion estimation system processor 14 will initially rectify the images as received from the camera 13. In the above description, the images I and I'are images as rectified by the ego-motion estimation system processor 14. To rectify the images as received from the camera 13, the camera 13 will need to be calibrated during a calibration operation prior to use in connection with recording images for use in estimating vehicle 10 motion as described above. Before describing operations to be performed during calibration, it would be helpful to consider the effects of incorrect calibration.

If, for example, a vehicle is driving along a straight road, and if the optical axis of the camera is aligned with the direction of motion, the flow field of successively recorded images will be an expansion field with the focus of expansion located at the center of the respective images, that is, at the origin (x, y) = (0, 0) of the image's rectilinear coordinate system. On the other hand, if the camera is mounted on the vehicle with a small rotation around the vertical ("Y") axis in three-dimensional space, then the focus of expansion will be displaced along the image's horizontal ("x") axis. In that case, the motion model defined by equation (11) will not account for the flow field, but will be well approximated by a forward translation and a rotational velocity wy around the vertical ("Y") axis.

Accordingly, errors in the orientation of the camera around the vertical ("Y") axis in three dimensional space will create a bias in the rotation estimate, in which case a curved path would be estimated instead of a straight path. Similarly, errors in the camera's orientation around the horizontal ("X") axis in three dimensional space will cause a bias in the pitch estimate. Based on these observations, a calibration operation can be performed by having the camera record a sequence of images while the vehicle is being driven down a straight roadway. The vehicle's ego-motion is estimated as described above in connection with FIGS. 2 or 3, and calibration parameters are estimated that would cause the ego-motion to integrate into a straight path.

The invention provides a number of advantages. In particular, the invention provides an arrangement for determining ego-motion of a vehicle 10 on a roadway from a series of images recorded by a camera 13 mounted on the vehicle 10, at least a portion of the images comprising projections of the roadway, and without requiring mechanical

Tesla EX1002 Page 1477 IPR2025-00943

sensors which are normally not provided with a vehicle 10 and that would, if provided, increase the cost and maintenance expenses thereof.

It will be appreciated that a system in accordance with the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof, any portion of which may be controlled by a suitable program. Any program may in whole or in part comprise part of or be stored on the system in a conventional manner, or it may in whole or in part be provided in to the system over a network or other mechanism for transferring information in a conventional manner. In addition, it will be appreciated that the system may be operated and/or otherwise controlled by means of information provided by an operator using operator input elements (not shown) which may be connected directly to the system or which may transfer the information to the system over a network or other mechanism for transferring information in a conventional manner.

The foregoing description has been limited to a specific embodiment ofthis invention. It will be apparent, however, that various variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention. It is the object of the appended claims to cover these and such other variations and modifications as come within the true spirit and scope of the invention.

What is claimed as new and desired to be secured by Letters Patent of the United States is :

Espacenet

Claims: EP1257971 (A4) — 2005-07-06

SYSTEM AND METHOD FOR ESTIMATING EGO-MOTION OF A MOVING VEHICLE USING SUCCESSIVE IMAGES RECORDED ALONG THE VEHICLE'S PATH OF MOTION

Claims not available for EP1257971 (A4)
Claims of corresponding document: WO0139120 (A2)

A high quality text as facsimile in your desired language may be available amongst the following family members:

CA2392652 (A1) WO0139120 (A2)

- Original claims
- Claims tree

The EPO does not accept any responsibility for the accuracy of data and information originating from other authorities than the EPO; in particular, the EPO does not guarantee that they are complete, up-to-date or fit for specific purposes.

CLAIMS 1. An ego-motion determination system for generating an estimate as to the ego-motion of a vehicle moving along a roadway comprising :
A. an image information receiver configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway ; and
B. a processor configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle.

European Patent Office

SUPPLEMENTARY EUROPEAN SEARCH REPORT

Application Number
EP 00 98 0706

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	EP 0 390 051 A (HONEYWELL INC) 3 October 1990 (1990-10-03) * abstract; figure 1 * * page 1, line 10 - line 14 *	1	G06T7/20 G05D1/02
X	EP 0 390 050 A (HONEYWELL INC) 3 October 1990 (1990-10-03) * page 12, line 21 - line 38; figure 7 *	1	
X	EP 0 786 739 A (NEC CORPORATION) 30 July 1997 (1997-07-30) * page 1, line 1 - line 52 *	1	
X	US 5 259 040 A (HANNA ET AL) 2 November 1993 (1993-11-02) * abstract; figures 1,2 * * column 4, line 33 - column 5, line 26 *	1	
X	US 5 257 209 A (MARKANDEY ET AL) 26 October 1993 (1993-10-26) * abstract; figures 1-4 * * column 3, line 1 - line 19 *	1	

TECHNICAL FIELDS SEARCHED (Int.Cl.7)

G08G
H04N
G06T
B60K
G05D
G01S

The supplementary search report has been based on the last set of claims valid and available at the start of the search.

Place of search	Date of completion of the search	Examiner
Munich	10 May 2005	Schinnerl, A

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

1

EPO FORM 1503 03.82 (P04C04)

IPR2025-00943
Tesla EX1002 Page 1480

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 00 98 0706

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

10-05-2005

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0390051 A	03-10-1990	US 4969036 A	06-11-1990
		EP 0390051 A2	03-10-1990
		JP 3048977 A	01-03-1991
EP 0390050 A	03-10-1990	US 5036474 A	30-07-1991
		EP 0390050 A2	03-10-1990
		JP 3223971 A	02-10-1991
EP 0786739 A	30-07-1997	US 5751838 A	12-05-1998
		DE 69717176 D1	02-01-2003
		DE 69717176 T2	30-04-2003
		EP 0786739 A2	30-07-1997
		JP 9219814 A	19-08-1997
US 5259040 A	02-11-1993	DE 69231812 D1	07-06-2001
		DE 69231812 T2	15-11-2001
		EP 0606385 A1	20-07-1994
		JP 6511339 T	15-12-1994
		JP 3253962 B2	04-02-2002
		WO 9307585 A1	15-04-1993
US 5257209 A	26-10-1993	NONE	

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 622 072 B1

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
07.07.2010 Bulletin 2010/27

(51) Int Cl.:
G06N 3/00 *(2006.01)*

(21) Application number: **04291912.6**

(22) Date of filing: **27.07.2004**

(54) **An automated action-selection system and method and application thereof for training prediction machines and for driving the development of self-developing devices**

Ein automatisiertes System zur Aktionsanwahl, sowie das Verfahren und dessen Anwendung, um Prognosemaschinen auszubilden und die Entwicklung sich selbst entwickelnder Geräte zu unterstützen

Système de sélection d'actions, méthode et application associées, pour former des machines de prévision et pour conduire le développement de dispositifs capables de s'auto-développer

(84) Designated Contracting States:
DE FR GB

(43) Date of publication of application:
01.02.2006 Bulletin 2006/05

(73) Proprietor: **Sony France S.A.**
92110 Clichy La Garenne (FR)

(72) Inventors:
- **Kaplan, Frédéric**
  75005 Paris (FR)
- **Oudeyer, Pierre-Yves**
  75013 Paris (FR)

(74) Representative: **Thévenet, Jean-Bruno et al**
Cabinet Beau de Loménie
158, rue de l'Université
75340 Paris Cédex 07 (FR)

(56) References cited:
- **F. KAPLAN, P. OUDEYER: "Motivational principles for visual know-how development" PROCEEDINGS OF THE THIRD INTERNATIONAL WORKSHOP ON EPIGENETIC ROBOTICS, 5 August 2003 (2003-08-05), XP002329049 BOSTON - USA**
- **F. KAPLAN, P. OUDEYER: "Maximizing Learning Progress: An Internal Reward System for Development" EMBODIED ARTIFICIAL INTELLIGENCE: INTERNATIONAL SEMINAR, 11 July 2003 (2003-07-11), XP002329050 DAGSTUHL CASTLE - GERMANY**
- **SCHMIDHUBER J ED - INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS: "Curious model-building control systems" PROCEEDINGS OF THE INTERNATIONAL JOINT CONFERENCE ON NEURAL NETWORKS. SINGAPORE, NOV. 18 - 21, 1991, NEW YORK, IEEE, US, vol. VOL. 2, 18 November 1991 (1991-11-18), pages 1458-1463, XP010054430 ISBN: 0-7803-0227-3**
- **P. OUDEYER, F. KAPLAN: "Intelligent Adaptive Curiosity: a source of self-development" PROCEEDINGS OF THE FOURTH INTERNATIONAL WORKSHOP ON EPIGENETIC ROBOTICS, 27 August 2004 (2004-08-27), XP002329051 GENOA, ITALY**

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 1 622 072 B1

Printed by Jouve, 75001 PARIS (FR)

Description

Field of the Invention

[0001] The present invention relates to the optimization of series of actions or trials which are performed in order to learn. More particularly, the invention is directed to an automated method and system for selecting which trials or actions should be tried next in order to achieve efficient learning.

[0002] The invention also relates to application of this technique in methods for training prediction machines, and to control the progress of active learning, or to organize the behaviour of, a self-developing device or system (e.g. a robot). The invention further relates to systems for training prediction machines using this technique and to systems applying this technique for controlling active learning or behaviour-organization in a self-developing device or system. The invention yet further relates to prediction machines trained using this technique and to self-developing devices applying this technique.

Background of the Invention

[0003] There is a wide variety of applications in which it is necessary or useful to perform a series of actions on or affecting a setup or system S in order to acquire knowledge about the behaviour of the system S. The setup S can be a natural system or an artificial (man-made) system - use of the expression "setup" does not imply that the system has been actively organized or configured (although that could be the case). Often the desired "knowledge" consists in learning some relationship that holds between a set of the system's parameters. In general, it is desirable to minimize the number of actions that must be carried out.

[0004] In general, any setup or system S can be described in terms of its own current state and the state of an environmental context in which the system finds itself. The system's current state can be described in terms of values taken by a number of different parameters characterizing the system ("system-state parameters"). In a similar way, the system's environmental context can be described in terms of values taken by a number of different parameters characterizing that context ("context parameters"). In many cases it is not possible to determine clearly whether a given parameter is a system-state or context parameter. Indeed this distinction is immaterial as far as the operation of the present invention is concerned. According both kinds of parameter will be referred to herein as system-state/context parameters.

[0005] Actions can be taken which affect the system S. For example, it may be possible to control a parameter which affects the system S, so that this parameter takes a selected value. The controlled (or "tuned") parameter may be a system-state/context parameter but this is not obligatory. For a particular system *S* it is possible to select a set of parameters (or "action variables") whose values will be tuned in actions that are intended to elicit information about the behaviour of the system *S*.

[0006] The present invention concerns the selection of actions, affecting a system *S*, which are to be taken deliberately with a view to observing how the system reacts or behaves. In the present document, the expression "action" is not intended to imply a necessity for physical motion or the operation of an actuator; on the contrary, the expression "action" relates to the deliberate setting of a set of one or more action variables to respective particular values. In many cases the tuned variables will relate to a physical quantity (e.g. the amount of water to supply to a crop, the level of a voltage to apply to an electrode, etc.), but this is not a requirement of the present invention.

[0007] It is helpful to explain what is meant in the present document by the expression "system-state/context/action space".

[0008] Consider an application in which a user wishes to discover the relationships inherent in a setup *S* in which circuit boards are being photographed for quality control purposes. The circuit boards can have substrates made of different materials and, hence, can have different brightness values. The user wishes to discover how the exposure time and ambient illumination affect the contrast between the circuit boards and the conductive paths they carry, as observed in the photographic images. More particularly, he wishes to find out what conditions produce high-contrast images.

[0009] It is assumed that the photographic equipment includes a digital camera mounted over a conveyor which carries the circuit boards into a fixed position relative to the camera. An image-processing unit determines automatically the contrast between the circuit boards and the conductive traces they carry, by processing the image data generated by the digital camera. For simplicity, it is assumed that the image processing unit can accurately determine what within the image corresponds to the circuit board and what corresponds to the conductive traces.

[0010] This particular system *S* can be described in terms of the brightness of the circuit board substrate, ***BR***, the ambient illumination, ***AI***, and the exposure period, ***EP***, at the time when the photographic images are generated. Only the exposure-period parameter ***EP*** can be set to different values under the user's control (or automatically), the other two parameters are considered to be outside the user's control. Thus the circuit board brightness, ***BR***, and ambient illumination, ***AI***, can be considered to be system-state/context parameters describing the system, and the exposure period ***EP*** can be considered to be an action variable whose value can be varied or "tuned".

IPR2025-00943
Tesla EX1002 Page 1483

[0011] It can be considered that these parameters ***AI***, ***BR*** and ***EP*** define a multi-dimensional space which, in this simple case, is a three-dimensional space as illustrated in Fig.1. In a more realistic example the multi-dimensional space is likely to have considerably more dimensions than three, however this would be difficult, if not impossible, to represent in a drawing.

[0012] The above-described multi-dimensional space defined by the system-state/context parameter(s) and action variable(s) of a system is referred to in the present document as "system-state/context/action space".

[0013] At any given moment, the above-mentioned example system can be described in terms of a vector defining the values taken by the circuit board brightness, ***BR***, ambient illumination, ***AI***, and the exposure duration, ***EP***. Fig.1 shows an example of a vector, ***A***, corresponding to one particular combination of values of ***BR***, ***AI*** and ***EP***.

[0014] The cuboid shown in Fig. 1 corresponds to the system-state/context/action space for the contrast-measurement system *S* described above. It will be noted that Fig.1 shows limits of 0 and 1 on the values of each system-state/context parameter and action variable. This represents the system-state/context/action space in the case where each of the system-state/context parameter values and action variable values is normalised so as to range from 0.0 to 1.0.

Description of the Prior Art

[0015] Various proposals have already been made in the field of statistics, and in the field of developmental robotics, with regard to how a series of trials or actions can be scheduled so as to optimize learning.

[0016] In the field of statistics, "optimal experiment design" seeks to determine how it is possible to minimize the number of examples that it is necessary to consider in order to achieve a given level of performance in generalization.

[0017] In the field of developmental robotics, one of the main goals is to produce robots which can develop. In this context "development" consists in a progressive increase in the complexity of the activities in which the robot can engage, with an associated increase in the robot's capabilities. As the robot develops, it can be said that the robot is engaged in "learning". In order to develop, the robot must explore its sensory-motor space; in other words, it must discover the consequences of performing given actions. For example, this could mean determining the visual changes which result from setting the speed of the robot's wheels to a particular value in a given environment (or from setting the joints in the robot's neck in a particular configuration, etc.). The way in which the robot explores the sensory-motor state-space is by performing a series of actions in a given context (or environment) and noting the perceived results. This set of interactions enables the robot to learn.

[0018] It is desirable to maximize the efficiency of robot learning, in other words to reduce the number of actions that the robot must perform in order to improve its capabilities to a particular extent. Randomly exploring the sensory-motor state-space would be immensely inefficient; to gain knowledge about its sensory-motor mapping, the robot needs to explore the sensory-motor state-space methodically.

[0019] In the field of developmental robotics, it has been found that the efficiency of robot learning can be improved by ensuring that the robot first tackles tasks that are relatively simple but then moves on to tackling tasks of progressively increasing difficulty. In general, humans control the situations or tasks encountered by the robot at a given time, ensuring that there is a progression towards tasks of increasing complexity. Thus the learning process can be described as "passive learning" or "passive development".

[0020] Currently, there is considerable interest in producing a mechanism to enable a robot (or other man-made system) to develop autonomously. In this case a learning process can be the engine of development. More particularly, it is desired to produce a robot (or other man-made system) which, when it encounters a complex, continuous, dynamic environment, is capable of determining, without pre-programmed knowledge, which tasks or situations in this environment have a complexity which is suited for efficient learning at a given stage in its development. A robot/man-made system of this type would learn easy parts of its sensory-motor mapping first of all and then progressively shift its focus of attention to situations of increasing complexity.

[0021] In the field of machine learning, various "active learning" proposals have been made as to how to choose actions which will provide a maximal gain in knowledge or "know-how". These proposals differ with regard to the manner in which "knowledge" and gains in knowledge are quantified. They differ, also, with regard to the way in which it is decided which action is liable to result in the maximal gain in knowledge. Some of the main known proposals are discussed below.

"Choose the action which gives the least predictable result"

[0022] One approach was proposed by J. Denzler and C. Brown in "Information Theoretic Sensor Data Selection for Active Object Recognition and State Estimation" (see IEEE Transactions on Pattern Analysis and Machine Intelligence, 24, vol. 2, pp. 145-157, (2001)). This paper concerns a robotic system which selects its next action based indirectly on the output of a prediction module which produces predictions of the sensory consequences of a given candidate action. For example, the prediction module might predict that a particular predicted variable will take a value 2.0. A decision-making module evaluates the degree of certainty of the predictions produced by the prediction module (e.g. "the predicted

IPR2025-00943
Tesla EX1002 Page 1484

value 2.0 will be correct 65% of the time) and selects, as the next action to perform, the candidate action for which the prediction module's prediction has the lowest degree of certainty.

[0023] In some applications this method can be efficient. However, in an environment which contains a random element, this approach can lead the robot to concentrate on taking actions which explore that random element (because it is maximally unpredictable). However, in general there is nothing useful to be learned by exploring a random element of this type. Thus, development of the robot will be side-tracked.

[0024] Moreover, techniques in which the certainties of predictions are evaluated are computationally expensive and cannot readily be implemented for continuous system-state/context/action spaces.

"Choose the action which is expected to result in the greatest prediction error"

[0025] Another approach was proposed by S. Thrun in "Exploration in Active Learning", in Handbooks of Brain Science and Neural Networks, ed. M. Arbib, MIT Press (1995). This proposal concerns another system which makes use of a prediction module which produces predictions of the consequences of a given candidate action. Once again, there is a module which evaluates the performance of the prediction module, this time by evaluating the prediction error that is expected to result from implementing a given candidate action. This decision-making module selects as the next action to be performed that candidate action for which it expects the greatest prediction error to arise.

[0026] This method can be implemented very efficiently and run in high-dimensional spaces in real-time on small robots. Moreover, it can be useful in controlled environments (e.g. when an artificial device is learning to navigate in a static environment). However, it is counter-productive in an uncontrolled non-deterministic environment because, once again, it can lead to a focus of attention on a random element in the environment.

"Choose the action which will maximally reinforce the certainty of the predictions which are currently uncertain"

[0027] A further technique was proposed by N. Roy and A. McCallum in "Towards Optimal Active Learning through Sampling Estimation of Error Reduction", in Proceedings of the 18th International Conference on Machine Learning (ICML-2001). This proposal arises in the statistical field of "optimal experiment design" and takes a probabilistic approach.

[0028] According to this technique, a robot contains a prediction module, arranged to predict the outcome of certain candidate actions. There is a further module which predicts the degree of certainty of these predictions. This further module is updated in a dynamic manner so that the reliability of the various predictions can become more certain. In this system, the candidate action which is selected for performance is the action which will lead to a maximum reinforcement of the certainty of the predictions for which the robot is currently unsure.

[0029] Unlike the above-described Denzler and Brown system, according to this technique the evaluation of a candidate action includes a step of predicting how the robot's knowledge will be changed, and how this will affect the certainties of its predictions. It will be seen, then, that in this system there is an attempt to evaluate the gain in knowledge that will arise from choice of a particular candidate action.

[0030] In an environment which contains a random element, this Roy and McCallum approach can avoid the situation where the robot is trapped in exploration of that random element. However, the robot needs to consider a certain minimum number of examples before it will determine that it is focusing on a random element (a part of the sensory-motor state-space or "system-state/context/action space" where its actions have no effect on the certainties of its predictions). Thus, time and computational resources are wasted on evaluating and performing actions which can lead to no useful gain in knowledge. This wasted time may be minimized in a system where the noise is independently normally distributed. However, that will not be the case in a non-controlled, realistic environment.

[0031] Furthermore, this method requires use of considerable computational resources and no efficient implementation exists for the case where the sensory-motor state-space (or, more generally, system-state/context/action space) is continuous.

"Choose the action which (you predict) will maximally decrease prediction error"

[0032] An earlier active learning proposal by the present inventors is described in copending European patent application no. 03 291 359.2. According to this technique, a self-developing device contains a prediction module designed to predict the outcome of candidate actions, and a meta-prediction module which predicts the prediction errors that are associated with the different predictions made by the prediction module. The self-developing device updates the meta prediction module, and monitors its current level of prediction errors. The candidate action that is selected for performance is that action which, the meta prediction module predicts, will lead to the largest decrease in prediction error (possibly several time steps in the future).

[0033] The above-described technique has proved useful in certain situations, and successfully prevents a robot from focussing its attention on a random element in the environment. However, in some cases this technique can lead the

IPR2025-00943
Tesla EX1002 Page 1485

robot to enter a loop in which it alternately focuses its attention on a first situation associated with a high prediction error rate, and a second situation associated with a low prediction error rate. For example, in the case of a robot moving in a space enclosed by walls, the robot would spend time alternately bumping randomly against walls (where the error rate of predicting what would happen next is high) and then stopping in front of the wall, looking at it (a situation in which it is simple to predict what will happen next).

[0034] The inventors' earlier paper "motivational principles for visual know-how development" (published in the Proceedings of the Third International Workshop on Epigenetic Robotics, 5 August 2003, Boston, USA) describes a self-developing robot which measures various parameters of its own state and its environment and attempts to determine, autonomously, which action to take next. This self-developing robot calculates a reward function associated with different possible actions and selects an action associated with a high value of the reward function. The reward function is based on the values of one or more particular task-independent motivational variables. One motivational variable is a "predictability" variable that is based on the accuracy with which the self-developing device can predict the outcome of a particular function. Calculation of the predictability motivational variable involves predicting what values will be taken by certain variables after a given action and then comparing the prediction with the actual values.

Summary of the Invention

[0035] In most systems, the behaviour of the system varies depending upon "where the system is" in system-state/context/action space. In other words, for different regions in the system-state/context/action space, different relationships exist between the various system-state/context parameters and/or action variables. Each different region in system-state/context/action space corresponds to a different type of situation in the system, ***S***.

[0036] Fig.2 illustrates the case where the system-state/context/action space of Fig.1 can be divided into three regions in which respective different relationships hold between the ambient illumination, substrate brightness and exposure period. These three regions are labelled R_1, R_{21} and R_{22}. The vector ***A***, corresponding to one particular combination of values of the system-state/context parameters and action variable, has an end-point in region R_1.

[0037] As mentioned above, when investigating the behaviour of a setup ***S*** and trying to select which action affecting the setup should be performed next, it can be helpful to consider what learning progress can be expected to result from the performance of different candidate actions. The present inventors have realised that the best results are not obtained when learning progress is assessed based on analysis of the immediately-preceding series of actions that have been performed. Instead, it is advantageous to consider the previous actions of the same kind as the present one that have been performed when the setup was in a similar situation to the present one.

[0038] Clearly, it is necessary to have some way of determining which actions are "similar" to the presently-considered one and which situations are "similar" to the current situation. According to the present invention, action/situation pairs are judged to be similar to each other if they occupy the same region in system-state/context/action space. In other words, a situation ***A*** corresponding to system-state/context parameter and action variable values $SCA_A(t)$ and a situation *B* corresponding to system-state/context parameter and action variable values $SCA_B(t)$ will be considered to be "similar" if $SCA_A(t)$ and $SCA_B(t)$ define vectors having end-points in a common region of the sensory-motor state-space.

[0039] The above concept is illustrated in a highly-simplified form in Fig.3 which shows three vectors, in the system-state/context/action space of Fig.1 having three regions defined as in Fig.2. In the system ***S***, photographs can be taken in the situations ***A***, ***B*** and ***C***, corresponding to combinations of values of the system-state/context parameters and the action variable that are represented by $SCA_A(t)$, $SCA_B(t)$ and $SCA_C(t)$, respectively. According to the present invention, the situations ***A*** and ***B*** are considered to be similar to each other because the respective vectors representing the combinations of values of the system-state/context parameters and action variable when the system is in these situations ***A*** and ***B*** have end-points in a common region (region R_1 in Fig.3).

[0040] On the other hand, situation ***C*** is not similar to either of situations ***A*** or B because the vector representing the corresponding combination of values of the system-state/context parameters and action variable when the system is in this situation ***C*** has an end-point in region R_{21} of Fig.3, which is different from the region R_1 containing the end-points of the vectors corresponding to situations ***A*** and ***B***.

[0041] According to the present invention, it is considered to be advantageous to use different functions for assessing learning progress for situations in respective different regions of system-state/context/action space.

[0042] According to a first aspect thereof, the present invention provides an automated action-selection method and system for selecting which action affecting a setup ***S*** should be performed next when seeking to acquire knowledge efficiently about the behaviour of the setup ***S***. By using the action-selection technique of the present invention, the number of actions that are required in order to gain knowledge about the behaviour of setup ***S*** can be reduced. Moreover, efficient learning can be achieved even in a complex, continuous, dynamic environment. Furthermore, the learning process does not become trapped in a loop between highly-predictable and highly-unpredictable situations, or become side-tracked into exploring a random element in the setup ***S***.

[0043] Preferred embodiments according to a first aspect of the present invention provide an automated action-se-

IPR2025-00943
Tesla EX1002 Page 1486

lection system and method which is capable of determining automatically which parts of system-state/context/action space should be explored next so as to match the complexity of the explored situation to the level of learning ("development") at the time.

[0044] The action-selection technique according to preferred embodiments of the first aspect of the present invention involves use of a prediction machine. More particularly, the prediction machine attempts to predict a respective value for each of a set of one or more variables, ***VAR***, that is dependent on or affected by the behaviour of setup ***S***. In order to select which action should be performed next, a set of possible next actions ("candidate actions") is evaluated. For each candidate action, the prediction machine predicts what respective value(s) it expects the predetermined set of variables ***VAR*** to take after this candidate action has been applied to the setup under investigation, taking into account the current state and context of that setup. A meta prediction machine then assesses how accurate it expects those predictions to be, notably by predicting what will be the prediction error.

[0045] One of the candidate actions is selected for performance so as to affect the setup ***S***. After a selected action has been performed, feedback is received regarding the effect of this action on the setup's state and context parameters and on the values of the set of predicted variables (unless these values can already be derived from the observed values of the system-state/context parameters). This feedback is used to adapt the function applied by the prediction machine, whereby to improve the accuracy of its predictions. This feedback is also used by the meta prediction machine so as to improve the accuracy of its assessments in the future.

[0046] According to the present invention, an assessment is made of the gain in knowledge that is expected to result from performance of a given candidate action. In certain preferred embodiments according to the first aspect of the invention, this expected gain in knowledge is evaluated by determining what reduction in prediction error is expected to result from performance of this candidate action.

[0047] A particularity of the present invention is that the expected gain in knowledge is not evaluated by comparing the level of knowledge immediately before the selected candidate action has been performed with the expected level of knowledge immediately after performance of this action. On the contrary, the expected gain in knowledge is evaluated by comparing the expected level of knowledge after the selected candidate action with the level of knowledge that existed on one or more previous occasions when the setup had values of system-state/context parameters and action variables that are similar to those applicable at the time when the selected candidate action will be performed.

[0048] In other words, according to the present invention learning progress is assessed based on study of some parameter which indicates the system's level of knowledge, and the evolution of this "knowledge-indicator" parameter is evaluated "locally" (that is, in the region of system-state/context/action space which contains the current values of the system-state/context parameter(s) and action variable(s)).

[0049] Because the expected change in the value of the knowledge-indicator parameter is evaluated based on actual values of this knowledge-indicator parameter observed in a comparable situation to the present one, and not simply based on the immediately preceding values of the knowledge-indicator parameter, the technique of the present invention enables efficient learning to take place whilst avoiding becoming trapped in a loop between highly-predictable and highly-unpredictable situations. Moreover, the technique of the present invention manages to avoid focussing attention on situations which contain a random element (and which, hence, cannot give rise to the learning of general rules).

[0050] In preferred embodiments of the invention, a candidate action is considered to produce a large knowledge gain if performance of this action is expected to cause a large decrease in prediction error. In other words, a prediction is made related to performance of a candidate action, and the expected prediction error is computed. The expected prediction error for this candidate action is compared with actual prediction error values observed when this kind of action was performed previously in similar circumstances to the present ones (that is, when the values of system-state/context parameter(s) and of the action variable(s) were in the same region of system-state/context/action space as they are now).

[0051] Once a candidate action is selected for performance, and has been performed, the data defining this action, the system-state/context parameters when this action was performed and the consequences of performance of this action (in terms of the observed values of the set of predicted variables, ***VAR***) can be considered to constitute a "training example".

[0052] Each training example can be represented using a multi-dimensional vector having as its different components the respective values of each system-state/context parameter at the time when this action was performed, the respective values of each action variable defining this action, and the respective values of each of the set of variables ***VAR*** that were observed as a result of performance of this action. According to this representation, each training example can be considered to be an association of two vectors, the first one defining the position of this training example in system-state/context/action space and the second defining the location of this training example in "outcome space" (that is, a space whose dimensions correspond to the actual values of the different predicted variables ***VAR***).

[0053] According to the preferred embodiments of the action-selection technique according to the first aspect of the invention, it is preferable to divide system-state/context/action space into regions in a progressive manner: starting by dividing the whole space up into a single region (or small number of regions) and then sub-dividing that region (or those

IPR2025-00943
Tesla EX1002 Page 1487

regions) later on.

[0054] As actions are selected and performed, training examples are accumulated. The training example data is stored in association with the region in system-state/context/action space which contains the end-point of the vector defining this training example. When a decision is taken to divide a region of system-state/context/action space into two or more new regions, the training examples that previously were associated with the original region become assigned to one or other of the new regions. Thus, at a given moment each training example is associated with only one region.

[0055] According to the preferred embodiments of the action-selection technique according to the first aspect of the invention, a given region of system-state/context/action space will be sub-divided into two (or more) smaller regions when a first criterion is met. According to one example of this first criterion, sub-division of a region, ***R***, will take place when it is determined that the latest training example is the Nth training example having a combination of system-state/context parameters and action variables which falls within this region R (with N being a predetermined number).

[0056] Preferably, when a region, ***R***, is sub-divided into two or more smaller new regions, the boundaries separating the new regions from each other will be set based on a second criterion. According to one example of this second criterion, the region boundaries are set so as to minimize the variance in "outcome space" of the locations of the training examples which are assigned to the same new region of system-state/context/action space (whilst keeping a balance between the numbers of training examples located in the new regions of system-state/context/action space).

[0057] It will be seen that as system-state/context/action space is sub-divided in this progressive manner a tree structure builds up: the initial region can be considered to be the root, the sub-divisions of this initial region constitute the branches, sub-divisions of those sub-divisions constitute twigs, etc. A region at the end of a given chain of sub-divisions can be considered to be a "leaf".

[0058] When an action is selected and performed, the action-selection system according to the preferred embodiment of the first aspect of the invention receives feedback enabling it to determine the actual prediction error that occurred. The actual prediction error data will be stored in relation to a particular region of system-state/context/action space, namely that region corresponding to the combination of values of system-state/context parameters and action variable(s) applicable at the time when the prediction was made. The size of this region will depend upon whether this combination of values is located in a part of system-state/context/action space that has been highly sub-divided or barely divided at all.

[0059] In the action-selection technique according to the first aspect of the invention, it has been found to be advantageous if, every so often, there is a random selection of the next action. In other words, on a majority of occasions the next action will be selected so as to maximise the expected knowledge gain (evaluated locally) but, on a minority of occasions, there will be a random selection of one of the candidate actions.

[0060] According to a second aspect, the present invention provides a system and method for training a prediction machine. The prediction machine is trained by causing it to participate in the selection of a series of actions using the action-selection technique according to the first aspect of the invention.

[0061] In the prediction-machine training method according to preferred embodiments of the second aspect of the present invention, for each of a set of candidate actions which can be taken to affect a setup ***S***, the prediction machine attempts to predict what would be the value of a set of variables, ***VAR***, if that candidate action were to be performed (the respective values of the set of variables ***VAR*** being dependent on or affected by the behaviour of the setup ***S*** and the setup ***S*** being describable in terms of a number of system-state/context parameters). A meta prediction module evaluates the expected accuracy of the prediction (expected prediction error), and a knowledge-gain assessment module assesses what reduction in prediction error would result from performance of the candidate action. The expected knowledge-gain is assessed by comparing an indicator of prediction error expected to occur if this action is performed with a comparable indicator of the actual prediction error that existed on one or more previous occasions when the setup had a combination of values of system-state/context parameters and action variables that is similar to that applicable at the time of the current experiment.

[0062] According to the prediction-machine training technique of the second aspect of the invention, system-state/context/action space is divided up into regions in a dynamic manner. The prediction, meta prediction and knowledge-gain assessment are made separately for each region and, as the series of actions progresses (the number of training examples increases), they become better and better adapted to suit the properties of their assigned region. For each candidate action, the expected knowledge gain is assessed using a prediction function, meta prediction module and knowledge-gain assessment module assigned to the region in system-state/context/action space corresponding to that candidate action and the current input values of the system-state/context parameters.

[0063] An action-selection module selects one of the candidate actions for performance. On a majority of occasions the selected candidate action is the one which is expected to produce the greatest reduction in prediction error (as assessed "locally" in system-state/context/action space). The selected action is performed, feedback is received enabling the actual prediction error to be determined, and the operation of the prediction machine modules assigned to the relevant region in system-state/context/action space is modified taking into account the received feedback.

[0064] Once the prediction machine has been trained, in a sequence of actions selected using the action-selection technique of the present invention, it can be considered to be programmed with the relationships inherent in the setup

IPR2025-00943
Tesla EX1002 Page 1488

S. Such a trained prediction machine has commercial value in its own right, in view of the fact that it can be used to predict how the setup S will behave for a given combination of values for the state and context parameters, and under the effect of different actions. Assuming that the setup S has parameters which are related to one another in a non-random fashion, this prediction will tend to be accurate because the prediction machine has already learned the relationships that hold between the system parameters.

[0065] Thus, the second aspect of the present invention also provides a prediction machine trained according to the above-described prediction-machine training technique.

[0066] The action-selection technique and prediction-machine training technique of the first and second aspects of the present invention can be implemented using purpose-designed hardware or using suitably-programmed general-purpose computers. In either case, it is necessary to configure the purpose-designed machine/programmed computer so that it knows what are the system-state/context parameters applicable to the setup S that is to be investigated, what actions can be taken affecting the setup S, notably what are the action variables which can be tuned (and what are the possible tuning values or ranges), and what variables(s) (***VAR***) should be predicted.

[0067] The necessary configuring can be performed by the manufacturer before supplying the purpose-designed machine, computer program or programmed general-purpose computer to the end-user. Alternatively, the purpose-designed machine, computer program, or programmed computer can be supplied to the end-user in a general state, allowing the user to identify what are the system-state/context parameters, action variables and predicted variable(s) (***VAR***) applicable to the setup he wishes to investigate. In the latter case, it may be possible for the user to re-initialise the purpose-designed machine/programmed computer so that he can input new system-state/context parameters, action variables and predicted variables in order to investigate a different setup.

[0068] After the initial configuration phase, a user, *U*, wishing to learn about the behaviour of a setup *S*, or to train the prediction machine to learn the behaviour of setup *S*, can input to the purpose-designed machine/programmed computer the current values of the system-state/context parameters. The purpose-designed machine/programmed computer will then indicate which action should be taken now. In suitable cases, the purpose-designed machine/programmed computer may automatically cause the selected action to be performed. Otherwise, the purpose-designed machine/programmed computer may simply indicate to the user which action should be taken next.

[0069] Once the selected action has been performed, the user inputs to the purpose-designed machine/programmed computer updated values of the system-state/context parameters and (if need be) the actual values of the set of predicted variables, and the purpose-designed machine/programmed computer will then determine what action should be performed next (what values should be assigned to the one or more action variables). Thus, the action-selection technique and prediction-machine training technique of the invention determine the sequence of actions that are performed affecting the setup S. This will enable the user or prediction machine to acquire knowledge about the setup S in a structured and efficient manner, reducing the number of actions that are required.

[0070] The action-selection technique and prediction-machine training technique according to the first and second aspects of the present invention are not limited to the case where a user (human or robot) inputs the current values of the system-state/context parameters and provides feedback on the consequences of performance of each action. In appropriate cases, the inputting of system-state/context parameter values and providing of feedback will be performed automatically, for example by sensors coupled to the purpose-designed hardware/programmed computer implementing the technique according to the invention.

[0071] According to a third aspect of the present invention, the above-described action-selection technique can be applied as a mechanism for driving the development of, or managing learning by, a robot or other self-developing device or system. More particularly, the third aspect of the present invention uses the action-selection technique of the first aspect of the invention to decide a sequence of actions to be performed by a self-developing device or system. The self-developing device or system performs the selected actions and provides feedback enabling the consequences of the action to be determined.

[0072] Preferred embodiments of the third aspect of the invention provide a decision-making module for a self-developing device or system. This decision-making module determines which action should be performed next by the self-developing device or system in order to promote efficient learning by that device or system and/or to achieve a desired organization of the behaviour thereof. The decision-making module is capable of determining for itself which parts of its system-state/context/action space should be explored next so as to match the complexity of the explored situation to the self-developing device or system's level of development at the time. The decision-making module can be integrated into the self-developing device or system itself, or can be physically separate from, but in communication with, that device or system.

[0073] The third aspect of the present invention enables active learning of complicated non-deterministic mappings in an efficient manner. Application of this approach by a robot (or other self-developing device or system) leads to organization of the robot's behaviour such that the robot selectively chooses the focus of its action (and "attention") in a manner leading to autonomous scaling of the complexity of the learning situations it encounters. The ability to organize the behaviour of a self-developing device or system in this way is an important achievement.

IPR2025-00943
Tesla EX1002 Page 1489

[0074] The method according to the third aspect of the present invention constitutes a drive motivating the selection of actions by a self-developing device or system. This drive can be termed "Intelligent Adaptive Curiosity". This drive can be termed "intelligent" because the action-selection method ensures that the self-developing device avoids situations that are overly predictable or overly unpredictable. This drive can be termed "adaptive" because the situations which will be attractive to the self-developing device change over time. This drive can be designated as "curiosity" because it maximizes learning progress, thereby pushing the self-developing device towards exploration of novel situations (where things can be learnt).

[0075] In many cases this "Intelligent Adaptive Curiosity" drive will be one of several drives motivating the behaviour of a self-developing device or system. For example, in a robot there may be a number of drives relating to maintenance of the robot's viability: for example, temperature maintenance, energy maintenance, maintenance of physical integrity, etc.

[0076] There are a wide variety of techniques allowing these different, and often competing, drives to be integrated with each other: for example, simple weighting of the different drives; allocation of reward functions to each drive and determination of the overall reward associated with a proposed action; etc. The technical literature in this field can be consulted for further examples. In a context where the method according to the third aspect of the present invention constitutes one of a number of drives motivating the behaviour of a self-developing device, the invention can be applied regardless of which technique is selected for integrating the different drives.

[0077] Further features and advantages of the various aspects of the present invention will become clear from the following description of preferred embodiments thereof, given by way of example, illustrated by the accompanying drawings, in which:

Fig.1 is a graph illustrating a multi-dimensional "system-state/context/action" space applicable to a simple setup or system, S;
Fig.2 is a graph illustrating different regions in the system-state/context/action space of Fig.1;
Fig.3 is a graph illustrating different vectors in the system-state/context/action space of Fig.1;
Fig.4 is a block diagram indicating schematically the main components of a preferred embodiment of action-selection system according to a first aspect of the present invention;
Fig.5 is a block diagram indicating schematically the main components of a preferred embodiment of prediction-machine training system according to a second aspect of the present invention;
Fig.6 is a schematic block diagram showing the main elements in one example of a fully-automated system applying the first (or second) aspect of the present invention in the field of biology;
Fig.7 is a schematic block diagram showing the main elements in one example of a fully-automated system applying the first (or second) aspect of the present invention in the field of finance;
Fig.8 is a schematic representation of an experimental set-up, involving a simple robot, used to test an embodiment according to the third aspect of the present invention;
Fig.9 is a block diagram indicating schematically the main components of a preferred embodiment of training system according to a third aspect of the present invention, for a self-developing device or system;
Fig.10 is a graph showing how prediction error evolves when actions are selected according to a preferred embodiment of the action-selection method of the present invention, by the simple robot in the experimental set-up of Fig.8;
Fig. 11 is a graph indicating the percentage of time the simple robot of Fig.8 spends in different situations when its actions are selected according to the preferred embodiment of the action-selection method of the present invention;
Fig. 12 is a graph showing how prediction error evolves when, according to a first comparative example, actions are selected in a random manner by the simple robot in the experimental set-up of Fig.8;
Fig. 13 is a graph showing how prediction error evolves when, according to a second comparative example, actions are selected according to a "maximizing" algorithm by the simple robot in the experimental set-up of Fig.8;
Fig. 14 is a graph indicating the percentage of time the simple robot of Fig.8 spends in different situations when actions are selected according to the "maximizing" algorithm;
Fig.15 is a graph showing how prediction error evolves when, according to a third comparative example, actions are selected according to a "minimizing" algorithm by the simple robot in the experimental set-up of Fig.8;
Fig.16 is a graph indicating the percentage of time the simple robot of Fig.8 spends in different situations when actions are selected according to the "minimizing" algorithm; and
Fig.17 is a graph comparing how learning progresses when the actions of the simple robot of Fig.8 are selected using the action-selection algorithm according to the preferred embodiments of the present invention and using algorithms according to the comparative examples.

FIRST ASPECT OF THE INVENTION

[0078] Fig.4 is a block diagram which illustrates schematically the main components of an action-selection system 1 according to a preferred embodiment of the first aspect of the present invention, which is used to investigate the behaviour

IPR2025-00943
Tesla EX1002 Page 1490

of a setup or system S. It is to be understood that, in practice, this action-selection system will often be implemented in software; the different "modules" shown in Fig.4 are identified merely to aid understanding of the various functions that are performed in the overall system. Moreover, the distribution of functions between the various modules shown in Fig. 4 could be changed and/or these functions could be performed using a lesser or greater number of modules than that shown in Fig.4.

[0079] As shown in Fig.4, the action-selection system 1 of the preferred embodiment of the first aspect of the invention includes a region-identifier module 2, a knowledge-gain assessment module 3, an action-selection module 4 and a module 5 for calculating actual prediction error. The action-selection system 1 selects which action should be performed at each of a series of time steps, these time steps being separated by a predetermined interval, τ.

[0080] The region-identifier module 2 receives signals SC(t) indicating what are the values of the system-state/context parameters of the setup or system ***S*** at a first time step t. This data can be obtained from the setup ***S*** in a number of different ways, as indicated above. In general, the values of the system-state/context parameters will either be input by a user or received directly from sensing devices coupled to the action-selection system 1.

[0081] The region-identifier module 2 also receives signals A(t) indicating what is the set of candidate actions that could be performed next. Typically, the kth candidate action, $A_k(t)$, assigns respective values v_1, v_2, ..., v_n to each of a set of n action variables AV_1, AV_2, ..., AV_n. This can be expressed, as follows:

$$A_k(t) = \{AV_1(t) = v_1; AV_2(t) = v_2; \ldots; AV_n(t) = v_n\}.$$

[0082] It will be noted that each candidate action assigns values to the action variables for the current time step ***t***. In fact, the action variables would not take the assigned values v_1 etc. until a time $t+\delta t$, because it takes the action-selection system 1 a finite time δt to choose a candidate action and to arrange for this candidate action to be performed. However, this will not perturb the good functioning of the action-selection system 1 provided that the system-state/context parameters do not change significantly from SC(t) during the period ***t*** to $t+\delta t$. In many applications this will inevitably be the case because the system's state and context will not change until the selected action is performed. In other applications, notably real-time applications, the time δt required to compute and select the action (typically a few milliseconds), and perform the action, is so short that the setup's state and context will not have changed during that brief time.

[0083] The set of candidate actions can be input to the action-selection system 1 of the present invention, for example by a human user or an external device or system. Alternatively, the action-selection system 1 may itself include a candidate action generator module (not shown). In either case, given the setup's current state and context, it may be impossible to assign certain values to the action variables. In order to avoid redundant processing, it is preferable for the set of candidate actions to be generated taking into account the current values of the system-state/context parameters, ***SC(t)***.

[0084] In certain preferred embodiments of the invention, the action variables and system-state/context parameters are normalized within the action-selection system 1 of the invention, so as to take values running from 0.0 to 1.0. This simplifies processing within the system 1. The actual inputs to and outputs from the system can be converted between the normalised values and the corresponding "real-world" value using well-known techniques. For example, the "real-world" range of a voltage variable, ***VV***, could be from 0 volts to 15 volts, such that the assignment of a value to this voltage variable in the following terms - "set ***VV*** to 0.2" - means "set voltage variable, ***VV***, to 3 volts".

[0085] The region-identifier module 2 has a dual role. Firstly, it serves to divide up system-state/context/action space into a number of regions, as will be explained in greater detail below. Secondly, for each candidate action, $A_k(t)$, the region-identifier module 2 identifies the region of system-state/context/action space that contains the combination of this candidate action $A_k(t)$ and the current values ***SC(t)*** of the system-state/context.

[0086] The action-selection unit 1 according to the invention is adapted to assess knowledge gain separately for the different regions of system-state/context/action space. This is achieved by arranging the knowledge-gain assessment module 3 to include a number of expert units, ***EXP,*** one for each region of system-state/context/action space, and by virtue of the fact that the region-identifier module 2 is capable of identifying which region of system-state/context/action space contains the combination of each candidate action ***A***(***t***) and the current values ***SC***(***t***) of the system-state/context parameters.

[0087] Fig.4 shows three expert units, $EXPR_{R1}$, EXP_{R21} and EXP_{R22}, in the knowledge-gain assessment module 3. However, the actual number of expert units defined in the knowledge-gain assessment module 3 at a particular time varies in a dynamic manner as the region identifier 2 sub-divides system-state/context/action space into new regions.

[0088] According to the preferred embodiment of the invention, each expert unit includes its own prediction module 12, meta prediction module 13, knowledge-gain assessment module 14 and data memory 15. When the region identifier 2 determines that the candidate action would put the setup ***S*** into a situation that is in a region R_{21} of system-state/context/action space, the expert unit EXP_{R21} associated with region R_{21} is used to predict the outcome of performing

IPR2025-00943
Tesla EX1002 Page 1491

this candidate action, to assess the expected knowledge gain that will result, and to output expected-knowledge-gain signals to the action selection module 4.

[0089] The operation of an expert unit EXP_{R21} associated with a region R_{21} of system-state/context/action space will now be considered in more detail. A counting variable q can be used to designate the successive occasions on which the setup S is in a situation corresponding to region R_{21} of system-state/context/action space. These occasions may be spaced apart in time in an irregular fashion.

[0090] When the region identifier 2 identifies region R_{21} as the region of system-state/context/action space that contains the situation corresponding to a given candidate action $A_k(t)$, taking into account the current values $SC(t)$ of the system-state/context parameters, the expert unit EXP_{R21} associated with region R_{21} is activated. For this candidate action $A_k(t)$, the prediction module 12 of expert unit EXP_{R21}, predicts the value of each a set of one or more variables VAR at the next time step (t+1). The nature of these variables VAR depends upon the system being investigated by the action-selection system of the invention. Typically, the prediction module 12 could attempt to predict what would be the value of one or more system-state/context parameters, or some variable(s) derivable from those parameters.

[0091] The predictions made by the prediction module 12 of expert unit EXP_{R21} can be labelled $VAR_{R21}(t+1)$, and they are based on the system's current system-state/context parameter values $SC(t)$, as well as on the given candidate action $A_k(t)$. In other words, $P_{R21}(A_k(t),SC(t)) = VAR_{R21}(t+1)$. More generally, for a candidate action $A(t)$ corresponding to a situation in region, R, of system-state/context/action space, the prediction can be expressed by the equation:

$$P_R(A(t),SC(t)) = VAR_R(t+1) \qquad (1)$$

[0092] A large number of different prediction techniques are known, for example: the nearest neighbour algorithm, methods using multi-layer perceptron neural networks, radial basis function neural networks, support vector machines, Bayesian machines, etc.. The present invention is not particularly limited with regard to which method is used to perform the prediction process. However, in preferred embodiments of the invention the prediction function applied by the prediction modules 12 of the expert units is the nearest neighbour algorithm.

[0093] Moreover, it is advantageous for the prediction implemented by the prediction module 12 of a given expert unit, for example EXP_{R21}, to be tailored to the characteristics of its region of system-state/context/action space, e.g. R_{21}. More particularly, over time the expert unit EXP_{R21} will adaptively modify the prediction function applied by its prediction module unit 12 so as to improve the prediction accuracy. (This will be the case for all of the above-listed examples of prediction techniques). The appropriate adaptation is implemented when the expert unit receives training example data for training examples located within the region of system-state/context/action space for which it is responsible. As each training example is assigned only to one region, this adaptation will be specific to the region R_{21}. Thus, even if the prediction modules 12 of all of the expert units initially apply the same prediction function (and initial use of the same function is not obligatory), after a time these functions will be different.

[0094] Each expert unit, EXP, also includes a meta prediction module 13 which estimates what will be the prediction error for each prediction made by the prediction module 12 of the same expert unit. For example, considering the operation of an expert unit EXP_R, it should be explained that, for each candidate action $A(t)$ corresponding to a situation in region R of system-state/context/action space, the meta prediction module 13 of this expert unit EXP_R computes a value indicating the respective errors $Ep_R(t)$ expected to be associated with the predictions $VAR_R(t+1)$ made by the prediction module 12 of this same expert unit, EXP_R, as follows:

$$MP_R(A(t),SC(t)) = Ep_R(t) \qquad (2)$$

[0095] Once again, various techniques (for example the nearest neighbour algorithm, methods using multi-layer perceptron neural networks, radial basis function neural networks, support vector machines, Bayesian machines, etc) are known for performing meta prediction and the present invention is not particularly limited with regard to which method is used to perform the meta prediction process. However, in preferred embodiments of the invention the meta prediction function applied by the meta prediction module 13 of expert unit EXP_R is a simple function according to which it is assumed that the prediction error for the currently-considered candidate action will be the same as the last actual prediction error observed for a training example in the region R for which this expert unit EXP_R is responsible.

[0096] Once again, the meta prediction function applied by the meta prediction module 13 of expert unit EXP_R can be tailored to the characteristics of the region R, such that the meta prediction function evolves over time whereby to increase in accuracy. Thus, even if the meta prediction modules 13 of all of the expert units initially apply the same prediction function (and, once again, this is not obligatory), after a time these functions will usually be different.

IPR2025-00943
Tesla EX1002 Page 1492

[0097] Each expert unit, ***EXP***, also includes a knowledge-gain assessment module 14 which examines the value of the prediction error that is expected to arise for a given candidate action, $A_k(t)$, corresponding to a situation in the region, ***R***, covered by this expert unit. The knowledge-gain assessment module 14 compares the expected prediction error with actual prediction errors that have been observed on one or more earlier occasions when the setup *S* had a combination of system-state/context parameters and action variables located in this same region R. This actual prediction error data is stored in a memory 15 associated with this expert unit. The operation of a preferred embodiment of the knowledge-gain assessment module 14 will be described in more detail later on.

[0098] The knowledge-gain assessment module 14 outputs, to the action selector module 4, data indicating the knowledge gain (decrease in prediction error) expected to arise if this particular candidate action is performed, together with data indicating the candidate action for which this expected knowledge-gain has been assessed.

[0099] On a majority of occasions, the action-selection module 4 selects, as the next action to be performed on or in the setup, *S*, that candidate action which is expected to lead to the greatest knowledge gain (which, in this case, is the greatest expected decrease in prediction error).

[0100] For a given set of candidate actions, it is fairly likely that knowledge-gain assessment modules 14 of more than one expert unit will output expected knowledge-gain information to the action selection module 4. In a variant of the present embodiment, the knowledge-gain assessment module 3 may itself include a unit which calculates which candidate action is expected to lead to the greatest knowledge gain. In this variant, the knowledge-gain assessment module 3 outputs to the action selection module 4 an indication of which candidate action is expected to yield the greatest knowledge gain (learning progress).

[0101] The action selection module 4 outputs a signal indicating what the selected candidate action is. This signal causes the selected action to be performed (either automatically or with intervention, for example, by a human user). The action selection module 4 also indicates to a calculator 5 which candidate action was selected for performance. The calculator 5 also receives a signal indicating the region, ***R***, of system-state/context/action space corresponding to the selected candidate action. This enables the calculator to retrieve the prediction $\boldsymbol{VAR_R(t+1)}$ from the appropriate prediction module within the knowledge-gain assessment module 3.

[0102] After an action is selected by the action-selection module 4 and performed at time t (more precisely, time $t+\delta t$), the calculator 5 determines what is the actual value, $\boldsymbol{VAR_{Ract}(t+1)}$, of the predicted parameter(s) at the time $t+1$. It may be that the respective values of the predicted variables in the set ***VAR*** can be calculated from the values of the system-state/context parameters. In such a case, the calculator 5 receives as inputs the values of the system-state/context parameters at the time $t+1$. Alternatively, it may be necessary or convenient to input to the calculator 5 the actual values of the predicted variables in the set ***VAR*** that are observed at time $t+1$ as the outcome of performing the selected candidate action.

[0103] Based on the inputs it receives, the calculator 5 calculates what is the actual prediction error, $\boldsymbol{E_{Ract}(t)}$, that is, what is the difference between the set of predictions $\boldsymbol{VAR_R(t+1)}$ made by the prediction module 12 for this candidate action and the actual outcome $\boldsymbol{VAR_{Ract}(t+1)}$ that has been observed.

[0104] For each predicted variable in the set ***VAR***, the actual prediction error, $\boldsymbol{e_{Ract}(t)}$, can be determined, as follows:

$$e_{Ract}(t), = [\text{absolute value } \{VAR_R(t+1) - VAR_{Ract}(t+1)\}]^k \qquad (3)$$

where k is an integer greater than or equal to 1 (typically, equal to 2). The overall prediction error $\boldsymbol{E_{Ract}(t)}$ can be calculated by combining the actual prediction errors obtained for each predicted variable (for example, by summing $|VAR_R(t+1) - VAR_{Ract}(t+1)|^2$).

[0105] The calculator 5 supplies this overall actual prediction error data $\boldsymbol{E_{Ract}(t)}$ to the memory 15 associated with expert unit $\boldsymbol{EXP_R}$, so that it can be taken into account by this expert unit in later calculations.

[0106] Now it is helpful to replace the variable t by a corresponding value of the counting variable q (recalling that q is the variable counting the number of times the system has encountered a situation in region R of system-state/action/context space, in other words q counts the number of training examples observed in region ***R***). If q^t indicates the value of counting variable q at time t, then the actual prediction error observed at time t can be designated $E_{Ract}(q^t)$.

[0107] At time step t the system is likely to have already been in situations in region R of system-state/context/action space on a number of occasions. Thus, there will be a series of actual prediction error values, $E_{Ract}(1)$, $\boldsymbol{E_{Ract}(2)}$, $\boldsymbol{E_{Ract}(3)}$, ..., $\boldsymbol{E_{Ract}(q^t-1)}$, stored in the memory 15 associated with expert unit $\boldsymbol{EXP_R}$. It can be impractical to base calculations on the entire series of actual prediction error values observed for the region ***R***, so it is preferred to consider only the $N+1$ last observed values for this region. Moreover, in order to reduce the effects of noise, the average is determined of the current value and the N previous values of actual prediction error, as follows:

IPR2025-00943
Tesla EX1002 Page 1493

$$Em_{Ract}(q^t) = \frac{1}{N+1}\{E_{Ract}(q^t) + E_{Ract}(q^t\text{-}1) + E_{Ract}(q^t\text{-}2) + \ldots + E_{Ract}(q^t\text{-}N)\} \quad (4)$$

It will be seen that the successive terms in the above equation, reading from right to left, correspond to successive occasions on which the system encountered situations in region ***R*** of system-state/context/action space. In general, these successive occasions will be spaced apart in time in an irregular manner.

[0108] Now, the knowledge-gain assessment module 14 of the expert unit EXP_R can define knowledge-gain (or "learning progress") achieved at time t, on the q^t-*th* occasion that the system 1 encounters a state in region ***R***, as follows:

$$LP_R(q^t) = Em_{Ract}(q^t) - Em_{Ract}(q^t\text{-}N) \quad (5)$$

[0109] Now, at the time when the knowledge-gain module 14 estimates the learning progress expected to be achieved because of performance of a given candidate action at time t, it does not yet know the actual prediction error that will result from performance of this action, so it cannot calculate the average value $Em_{Ract}(_qt)$. However, the knowledge-gain module 14 does have available to it the estimate $Ep_R(t)$ produced by the meta prediction module 13 of this same expert unit EXP_R. Accordingly, the knowledge-gain assessment module 14 uses this estimate $Ep_R(t)$ instead of $E_{Ract}(q^t)$ when calculating a value for $Em_{Ract}(q^t)$, and assesses expected learning progress by comparing the thus-calculated value of $Em_{Ract}(q^t)$ with $Em_{Ract}[(q^t\text{-}N)]$.

[0110] In the preferred embodiments of the first aspect of the present invention, the action selection module 4 is arranged so that on a majority of occasions it selects the candidate action for which the expected learning progress is the greatest. However, it is beneficial if, on some occasions, the action selection module 4 makes a random choice of the next action. This allows the action-selection system 1 to discover new regions in system-state/context/action space that are interesting from the point of view of yielding information as to the behaviour of the setup or system S.

[0111] In certain successful experiments (described below in relation to Figs. 10, 11 and 17), the action selection module 4 was arranged so that there was a probability of 0.85 that it would select the candidate action expected to yield the greatest learning progress, and a probability of 0.15 that it would select a random action. The present invention is not particularly limited with regard to the value of this latter probability. As long as the probability of selecting an action at random is greater than zero, the system 1 will discover new interesting regions of system-state/context/action space. However, if the value of this probability is equal to 1, then the system 1 will just be selecting actions at random. So, in the preferred embodiments of the invention, the probability, ***p***, that the action selection module 4 selects an action at random is $0<p<1$.

[0112] In some applications, the action selected by the action selector 4 will be automatically performed, under control of the action-selection system 1. In other cases, the action-selection system 1 may merely output signals indicating what are the desired values of the respective one or more action variables. In the latter case, the selected action is performed either by a human user or by another machine or system. The outcome of the action is fed back to the action-selection system 1, either in terms of the resulting values of the system-state/context parameters at some later time (e.g. *t*+1), or directly in terms of the actual values of the set of predicted variables ***VAR*** as observed at the later time (e.g. *t*+1). The action-selection system 1 associates this outcome data with data defining the performed action as well as with data defining the system description at the time when the action was performed. The resultant ensemble of data constitutes one training example.

[0113] As mentioned above, the training example data comprises:

- the performed action (that is, the values of the set of action variables defining the selected action performed at time t);
- the description of the setup/system S at time t (that is, the values of the set of system-state/context parameters at time t); and
- the outcome of the performed action (that is, the actual values of the set of predicted variables ***VAR*** at a predetermined later time, e.g. time *t*+1).

Incidentally, it is not essential for the feedback to be received by the system 1 at the time step subsequent to performance of the action. A number of time steps may intervene.

[0114] The training example data is supplied to the region identifier 2, as will be explained in greater detail below.

[0115] So far nothing has been said with regard to the way in which system-state/context/action space is divided up into the various regions. It is advantageous if this division of system-state/context/action space is performed by the

IPR2025-00943
Tesla EX1002 Page 1494

region-identifier module 2 itself, progressively refining regions into smaller and smaller areas as it learns more about the different situations that can present themselves. According to the preferred embodiment of the first aspect of the present invention, this can be achieved as follows:

[0116] The action-selection unit 1 is arranged to perform memory-based learning. In other words, the action-selection system 1 is adapted to "remember" (store data on) each of the training examples encountered since the system 1 was last initialised, that is to log training example data for each time step. This training example data is made available to the region identifier module 2, and is associated with the particular region of system-state/context/action state which contains the end-point of the vector defined by this training example's values of system-state/context parameters and action variables.

[0117] Initially the region identifier 2 has no information as to which situations are similar to each other; the whole of system-state/context/action space is considered to be a single region and there is a corresponding single expert unit in the knowledge-gain assessment module 3. As the action-selection system 1 operates it will build up a set of training examples corresponding to values (system-state/context parameter values, actual values of the predicted variables, etc.) encountered at each time step. The example data is allocated to the single existing expert unit which handles predictions for the single existing region of system-state/context/action space.

[0118] As data is gathered on more and more training examples, the region identifier 2 applies a criterion ***C1*** in order to decide whether or not to spit the initial single region of system-state/context/action space into two or more new regions. If the initial region is split into new regions, the original expert unit is split into a plurality of expert units - one for each of the new regions - (or new expert units are created), and the training example data that has already been collected is re-allocated to the expert unit associated with the relevant new region (that is, the new region which contains the combination of system-state/context parameter values and action variable values applicable at the time of this training example).

[0119] According to preferred embodiments of the first aspect of the invention, the criterion ***C1*** used for deciding whether or not to divide up a region is a simple judgement based on the number of examples encountered in the original region. Notably, when the number of training examples encountered for a given region of system-state/context/action space has exceeded a predetermined number, ***NS***, a decision is taken to partition the initial region into a plurality of new regions (for example, two new regions). In the experiments discussed below with reference to Figs. 10, 11 and 17, the predetermined number ***NS*** was 250.

[0120] At the time of splitting a region, a criterion ***C2*** determines how the original region is split, that is, what positions should be chosen for the boundaries of the new regions. In preferred embodiments according to the first aspect of the invention, the following criterion ***C2*** is used: when a region ***R*** is split into new regions (for example, two new regions $\boldsymbol{R}_1$ and $\boldsymbol{R}_2$) the location of the boundary is set so that the example data that had already been collected for region R is split in a balanced fashion between the new regions and with a minimal variance. That is, after the split each new region will have inherited a number of training examples which is similar to the number of training examples inherited by the other new regions, and the variance of the example data will be minimized. This variance can be minimized within each new region of system-state/context/action space, or within "outcome space" (that is, within a space defined by the actual observed values of the set of predicted variables, ***VAR***), or within a higher-dimensional system-state/context/action/outcome space defined by the values of the sets of system-state/context parameters, action variables and predicted variables. This condition can be expressed, as follows:

$$\text{minimize } [\{\text{card } (\boldsymbol{R}1) \text{ x variance } (\boldsymbol{R}1)\} + \{\text{card } (\boldsymbol{R}2) \text{ x variance } (\boldsymbol{R}2)\}]$$

[0121] According to the preferred embodiments of the first aspect of the present invention, when a region of system-state/context/action space is split in this way, the process of determining the location of the boundaries consists in first deciding which dimension (which parameter or variable, or combination thereof) to cut and then deciding the location of the (or each) cut point. The boundaries can be defined in simple terms, e.g. the first new region will contain examples in which a variable $\boldsymbol{V}_1$ takes a value 0 to 0.5, the second new region will contain examples in which variable $\boldsymbol{V}_1$ takes a value 0.51 to 1.0. The region identifier 2 maintains a list of cut dimensions and cut values.

[0122] The above-described region-splitting process is recursively iterated. When the criterion ***C1*** is met for any of the regions recognised by the region identifier 2, that region in its turn is split into new regions. After a while, a tree structure develops, with the initial region (and its expert unit) forming the root, the first level of division of the root region corresponds to a branch which in its turn may be divided into further branches, ending in regions (and expert units) which can be considered to represent "leaves" of the tree structure. This tree structure corresponds to a segmentation of the input space, with each expert unit being a specialist for its own part of the system-state/context/action space.

[0123] The action-selection unit 1 continues to operate and, at each time step, the region identifier 2 determines what "leaf" region contains the combination of a given candidate action and the current values of system-state/context parameters. This determination is performed by comparing the candidate action variable values and the current values of

IPR2025-00943
Tesla EX1002 Page 1495

the system-state/context parameters with the cut data defining the various regions of the tree structure. The expert unit applicable for this leaf region then makes its predictions. When the expected-knowledge gain has been assessed for each of the candidate actions, one candidate action is selected for performance. When the actual result is available, data indicative of the actual values of the set of predicted variables (***VAR***) is stored, together with other data for this training example, in the memory of the applicable leaf expert unit.

[0124] When the action-selection system 1 is used to select a series of actions, after some time it will build up a decision-tree architecture which reflects substantially all of the relationships that exist between the system-state/context parameters and action variables of the system, *S*. There will be a levelling-off in the prediction errors observed when predictions are made by the various prediction modules 12 in the expert units. This evolution of the prediction error values can be used as a sign that the learning process has succeeded in acquiring the desired knowledge about the setup S. Accordingly, the series of actions can be stopped. The action-selection system of the invention can include a control module (not shown) adapted to monitor the evolution of the prediction error values in the different expert units whereby to trigger a halt in the operation of the action-selection system 1 when a stop criterion is satisfied (e.g. when the percentage change in actual prediction error from one training example to the next has fallen below a threshold value for every "leaf" region defined by the region identifier 2).

[0125] By using the action-selection mechanism described above, the preferred embodiment of the first aspect of the present invention avoids transition effects: that is, when switching from an unpredictable situation to a highly predictable situation the assessment of expected knowledge gain will be based upon the previous actions in the unpredictable situation not merely upon the previous actions in time.

[0126] Moreover, by progressively refining the definitions of the regions it recognises in system-state/context/action space, the action-selection unit of the preferred embodiments of the present invention ensures that the learning process will first master the easy parts of the mapping to be learned, then progressively move on to the more difficult parts of the mapping. Not all parts of system-state/context/action space will be explored in detail, only those parts which are determined to be interesting.

[0127] By combining region-identification based on a "decision tree" architecture with predictions based on the nearest neighbour algorithm the action-selection unit of the preferred embodiment of the invention is an extremely efficient learning approach. Learning is maximally fast and prediction is accurate. In particular, the prediction is extremely fast compared with plain nearest-neighbour prediction: for example, with one million examples and ***NS***=250, a plain nearest neighbours algorithm would involve one million computations, whereas the technique of the preferred embodiment of the first aspect of the present invention would involve only 250 computations (i.e. it is 4000 times faster). Moreover the technique according to the preferred embodiment of the first aspect of the invention gives more accurate results than plain nearest-neighbour prediction because the region-splitting reduces the over-fitting which characterizes the nearest-neighbours prediction algorithm.

THE SECOND ASPECT OF THE INVENTION

[0128] When the action-selection system 1 of the first aspect of the present invention has been run so as to select a series of actions, and the prediction errors in the various prediction modules have all reached plateau values, it can be considered that the action-selection system 1 has become an expert on the system or setup S. More particularly, the prediction modules of the various expert units will tend to produce fairly accurate predictions of the effect of different candidate actions in different setup conditions (that is in for different values of the system-state/context parameters).

[0129] In these circumstances, the combination of the region identifier 2 and the prediction modules 12 in the knowledge-gain assessment module 3 can be considered to constitute a prediction machine that has been trained to make accurate predictions with regard to the behaviour of the setup S. More particularly, the combination of the region identifier 2 and the prediction modules 12 can predict the respective values of the set of one or more variables, ***VAR***, for different combinations of values of the action variables and system-state/context parameters.

[0130] Thus, according to a second aspect thereof, the present invention provides a method of training a prediction machine by causing it to form part of an action-selection system according to the first aspect of the invention and causing that system to select a series of actions until prediction errors level off in the prediction modules. The second aspect of the invention also provides a prediction-machine training system implementing this method.

[0131] Fig.5 shows, schematically, the main components of a preferred embodiment of prediction-machine training system 11 according to a preferred embodiment of this second aspect of the invention. It will be noted that this prediction-machine training system 11 has a similar construction and functioning to that of the action-selection system 1 that is illustrated schematically in Fig.4. Accordingly, the same reference numerals are used in Fig.5 to designate elements having a similar function to those in the action-selection system 1 of Fig.4.

[0132] Once again, it is to be understood that this prediction-machine training system 11 would probably be implemented in software and the various "modules" shown in Fig.5 are provided for illustrative purposes only. The various functions performed in the prediction-machine training system 11 may be split between modules in different ways from

IPR2025-00943
Tesla EX1002 Page 1496

that shown in Fig.5 and/or a greater or lesser number of modules can be used to perform these functions.

[0133] The prediction machine training system 11 shown in Fig.5 includes a region identifier 2, knowledge-gain assessment module 3, action selector 4 and prediction error calculator 5 comparable to those used in the experiment-selection system 1 of Fig.4. As indicated above, it is the combination of the region identifier 2 and the prediction modules 12 of the knowledge-gain assessment module 3 which constitutes the prediction machine to be trained by this training system 11.

[0134] The prediction machine training system 11 operates in two modes and has an operation mode setting module 16 which controls the different components of the system 11 dependent upon the mode in which the system is currently operating. In the first mode of operation the prediction machine is being trained and, in the second mode of operation, the prediction machine is already trained and is making predictions.

[0135] When the prediction-machine training system 11 is in the first operating mode, the region identifier 2, knowledge-gain assessment module 3 (and its expert units), action selector 4 and calculator 5 operate substantially as described above in relation to Fig.4. In other words, a sequence of actions affecting a setup S is selected and performed, and training example data is accumulated regarding each performed action, the setup's description (state and context) at the time the action was performed and the outcome of the action. This training example data is used by the region identifier 2 to progressively build-up a decision-tree architecture based on dividing system-state/context/action space into finer and finer regions, corresponding to the different types of situation inherent in the setup *S*.

[0136] The training example data for each region is also used to adapt the prediction function applied by the prediction module 12 of the expert unit assigned to handle that region in the knowledge-gain assessment module 3. In other words, as the sequence of actions is selected and performed, the functions applied by the different prediction modules 12 in the knowledge-gain assessment module 3 become more and more closely adapted to the properties/behaviour of a respective region in system-state/context/action space that corresponds to a certain type of situation in the setup S.

[0137] The prediction modules 12 can be considered to be trained when the actual prediction errors made by the prediction modules 12 stop falling (that is, they reach plateau values). This can be detected by a module (not shown) that monitors the evolution of the actual prediction errors for each prediction module. When this situation is detected, the operation mode setter 16 switches over from the first operating mode to the second operating mode of the prediction-machine training system 11.

[0138] In the second mode, the operation mode setter 16 sends signals to the knowledge-gain assessment module 3 and action selector 4 so as to disable certain of their functions. More specifically, when the prediction-machine training system 11 is in the second mode of operation, the user inputs data identifying a possible action that could be performed, that is, specifying values for the set of action variables. The user also inputs data describing the setup's state and context, more particularly data specifying respective values for the set of one or more system-state/context parameters.

[0139] The region identifier 2 identifies which region of system-state/context/action space contains the end-point of the vector defined by this set of action variable and system-state/context parameter values, and activates the prediction module 12 of the expert unit which handles this region. The activated prediction module outputs its prediction of the values of the set of one or more predicted variables, VAR, to a prediction-outputting interface 18, which supplies them to the user.

[0140] In some applications, the user may use the trained prediction machine to make a prediction of the outcome of a particular action, perform that action, and feed back to the system 11 the outcome of the performed action. In such a case, the fed-back data can be used to generate a further training example which can be used to further adapt the prediction function applied by the appropriate prediction module.

EXAMPLES OF APPLICATIONS OF THE FIRST AND SECOND ASPECTS OF THE PRESENT INVENTION IN DIFFERENT FIELDS

[0141] The action-selection/prediction-machine-training technique of the present invention can be applied in a wide variety of fields. A few examples of applications in different fields will be outlined below. In each of the examples below it is assumed that an action-selection system (prediction-machine training system) as described above in relation to Fig. 4 (Fig.5) is used.

Example 1 - Biological Field

[0142] In this particular example, a human user wishes to find out how a certain organ of an organism reacts when injected with a mixture of chemicals. The mixture of chemicals can be defined by specifying the concentration of each chemical ($C_1, C_2, C_3, ..., C_n$) and the reaction of the studied organ can be measured by determining what is the concentration (CC_{30}) of cells of a certain type in that organ thirty minutes after injection of the mixture of chemicals.

[0143] Each trial or action involves injection of chemicals into the organ and a consequent reaction which could entail damage to or destruction of that organ. Time and costs are also involved in performing each trial. Clearly it is desirable

IPR2025-00943
Tesla EX1002 Page 1497

to be able to minimize the number of trials that are performed.

[0144] The user can ensure that knowledge regarding the reaction of the organism to chemical mixtures can be acquired in an efficient manner, thereby reducing the number of trials that are required, by making use of an action-selection system according to a preferred embodiment of the present invention, such as that shown in Fig.4.

[0145] Preferably, the user employs a fully-automated assembly comprising the action-selection system of the present invention, together with a set of one or more automated robots which performs the selected actions and provides feedback to the action-selection system.

[0146] An example of such a fully-automated system is shown, schematically, in Fig.6. In the case illustrated in Fig. 6, an action-selection system according to the first aspect of the present invention (e.g. as shown in Fig.4) is used to select between different candidate actions that are possible at a given time. The selected action is output to an automatic laboratory which includes a set of experimenter robots. The selected action is performed automatically and feedback on the outcome is fed from the automatic laboratory back to the action-selection system.

[0147] More particularly, in this example the action-selection system 1 may be configured to consider the concentration values C_1 to C_n as action variables. In this example, the action-selection system will be configured such that its prediction module(s) seek to predict what CC_{30} will be, that is what will be the concentration of cells of the specified type in the organ thirty minutes after injection of a particular mixture of chemicals.

[0148] In this particular example, there are no system-state/context parameters to be taken into consideration. Thus, in this case the region identifier 2 and the actual prediction error calculator 5 do not have system-state/context parameter inputs from the setup S under investigation, the only input they receive from the setup S is the actual value of CC_{30} at a given time. However, if desired, the user could have configured the action-selection system 1 so that it recognises as system-state/context parameters (and monitors) the temperature, blood pressure etc. of the organism containing the organ under test.

[0149] The relevant prediction modules 12 of the action-selection system predict CC_{30} for each of a number of candidate actions, that is, for each of a number of different possible combinations of values of the action variables ***C1***, ***C2***, ..., ***Cn***. In the fully-automated assembly, the different sets of candidate actions that are considered may be generated automatically within the action-selection system 1 itself, or by a separate candidate-action generation module, or by the human user.

[0150] The candidate action to be taken is chosen (on a majority of occasions) based on the result of an evaluation by the action selector 4 as to which candidate action is expected to yield the greatest increase in knowledge compared to the last occasion on which the system had a combination of values of C_1 to C_n that was in the same region of system-state/context/action space as the combination specified in the candidate action. On a minority of occasions, a random action is selected. Output is given as to which mixture of chemicals has been selected. This output is supplied to the set of automated robots and these mix up the selected mixture of chemicals and inject it into the organ under investigation.

[0151] After thirty minutes, one of the automated robots measures the concentration (CC_{30}) of the specified cells in the organ. The automated robot supplies the measured concentration (CC_{30}) to the region identifier 2 and actual prediction error calculator 5 of the action-selection system 1 as feedback, enabling the system 1 to determine what was the actual error on the prediction made by the prediction module for this candidate action (this mixture of chemicals).

[0152] Initially the action-selection system 1 will recognise only a single region in system-state/context/action space. However, as more and more actions are performed, new expert units will be created and a decision-tree structure will develop sub-dividing the system-state/context/action space (which, in this case, reduces to "action space") into different regions. The structure of this decision tree will reflect the relationships that hold between the concentrations of the various chemicals injected into the organ and the reaction of the organ to those chemicals (to the extent that any such relationships exist).

[0153] When each relationship has been learned by a prediction module 12 of the action-selection system 1 according to the invention operating on the relevant region of system-state/context/action space, the prediction errors made by the respective prediction module 12 will level off at a plateau value. The action-selection system 1, or the user, can monitor the evolution of the prediction error values so as to determine when the prediction modules 12 have, in effect, learned all that they can. At this point the series of actions can be stopped.

[0154] At the end of the series of actions the action-selection system 1 of the invention has developed into a prediction machine that is trained to predict how the studied organ will react to different mixtures of chemicals. If the user is interested in knowing what will be the effect of administering a particular mixture of chemicals to the studied organ, he can simply input the appropriate combination of values of C_1 to C_n to the action-selection system 1, as if it were the only candidate action, and then sample the output from the appropriate prediction module - this will be a prediction of the resultant CC_{30}.

[0155] As an alternative to using the fully-automated assembly described above, the user could make use only of the action-selection system 1 of the invention and himself mix-up and administer the selected combination of chemicals and provide feedback to the action-selection system. The human user may also generate the various sets of candidate actions to be evaluated by the action-selection system 1 (although this can be done within the action-selection system 1 itself or by a separate candidate-action generation module).

IPR2025-00943
Tesla EX1002 Page 1498

[0156] It will be seen that, in this example of an application in the field of biology, the first aspect of the invention has provided the user with a machine which selects the actions in a sequence which enable the user to learn the effect of injecting particular combinations of drugs into an organ. The first aspect of the invention also provides a fully-automated assembly in which the action-selection system 1 determines which actions should be performed and a set of one or more automated robots performs the selected actions and provides feedback on the outcomes.

[0157] Moreover, in this example of an application in the field of biology, if the action-selection system of the invention were to be replaced by a prediction-machine training system such as that shown in Fig.5, the second aspect of the invention would provide the user with a prediction-machine training system which, in an interactive process, trains a prediction machine so that it can predict the effect of injecting particular combinations of compounds into an organ. As a result of this interactive process, the user is also provided with a prediction machine that is trained to make predictions of what will be the concentration of cells of a particular type in an organ 30 minutes after injection of different mixtures of compounds. This trained prediction machine could, for example, take the form of a computer system or an expert system.

Example 2 - Agricultural Field

[0158] Next consider an example in which a biotechnology firm is designing a particular species of yeast intended for sale to farmers. Yeast can be cultivated under different conditions, for example, using a greater or lesser quantity of water, fertilising chemicals and pesticides. The quality and quantity of the adult yeast will vary depending on the settings of these parameters as well as based on the agricultural context (for example, the type of soil, amount of sun, etc.). Different farmers may have different priorities: one wishing to use few chemicals whilst still assuring sufficient yield, another wishing to maximise the yield of yeast even if large quantities of chemicals are required.

[0159] The biotechnology firm wishes to study the way in which the amounts of water, fertilising chemicals and pesticides used during cultivation of its yeast affect the final yield, in different agricultural contexts.

[0160] The quantity of water, fertilising chemicals and pesticide applied per week during cultivation of the yeast can be specified using parameters ***W***, ***F*** and ***P***, respectively, the soil quality can be specified using a number of parameters $\boldsymbol{SQ_1}$, $\boldsymbol{SQ_2}$, etc., and the yield can be specified using a parameter ***Y*** indicating the number of kilograms produced per hectare a given number of weeks after planting.

[0161] Each action or trial involves cultivating a crop of the biotechnology company's yeast using water, fertilising chemicals and pesticide. There are labour costs, costs of water/chemicals, and extended periods of time required for performing each trial. Thus, it is desirable for the biotechnology company to be able to minimize the number of trials that are performed.

[0162] In this example, the action-selection system according to a preferred embodiment of the invention, such as that of Fig.4, may be configured to recognise the parameters ***W***, ***F*** and ***P*** as action variables, and to input as system-state/context parameters the values SQ_i defining the nature of the soil (e.g. in terms of its nitrogen content, clay content, etc.).

[0163] In this example, the action-selection system 1 will be configured such that its prediction module(s) 12 seek to predict what will be the yield ***Y*** of yeast at the end of a specified time period from planting.

[0164] The action-selection system 1 according to the preferred embodiment of the invention will predict ***Y*** for each of a number of candidate actions, that is, for each of a number of different possible combinations of applied quantities of water, fertilising chemicals and pesticides (different combinations of values of action variables ***W***, ***F***, ***P***), taking into account the current values of the system-state/context parameters $\boldsymbol{SQ_i}$.

[0165] Once again, on a majority of occasions the candidate action that is selected by the action selector 4 for performance will be the one that is expected to produce the greatest increase in knowledge compared to the last occasion on which there was a combination of values for **W**, ***F***, ***P*** and $\boldsymbol{SQ_i}$ that was in the same region of system-state/context/action space as the combination of current $\boldsymbol{SQ_i}$ values with the ***W***, ***F*** and ***P*** values specified in the candidate action. On a minority of occasions a random action will be selected. Output is given as to which combination of values of ***W***, ***F***, ***P*** has been selected.

[0166] The biotechnology company receives the output from the action-selection system 1 and may then cultivate a crop of yeast under the specified conditions. After the requisite time period, the yield ***Y*** is determined. The measured yield ***Y*** is input to the action-selection system 1 as feedback, enabling the system's calculator 5 to determine what was the actual error on the prediction made by the relevant prediction module 12 for this candidate action (this set of cultivation conditions), taking into account the values $\boldsymbol{SQ_i}$ at the time of the prediction.

[0167] As in the first example, as the series of trials progresses the action-selection system 1 will develop expert units which sub-divide system-state/context/action space, developing a decision-tree structure which reflects the relationships inherent in the system.

[0168] At the end of the series of trials, the combination of the region identifier 2 and the prediction modules 12 of the various expert units of the action-selection system 1 has developed into a prediction machine that is trained to predict

IPR2025-00943
Tesla EX1002 Page 1499

how the yield of the biotechnology company's yeast is affected by different cultivation conditions. This prediction machine can be used by the biotechnology company, or supplied to farmers along with the biotechnology company's yeast, so as to enable them to evaluate the yield likely to result from use of different combinations of cultivation conditions, given their soil conditions.

[0169] It will be seen that, in this example of an application in the field of agriculture, the first aspect of the invention has provided the biotechnology company with a machine which selects the actions in a sequence of actions which enables the company to learn the effect of different cultivation conditions on the yield of its new yeast.

[0170] Moreover, in this example of an application in the field of agriculture, the second aspect of the invention provides the biotechnology company with a prediction-machine training system which, in an interactive process, trains a prediction machine so that it can predict the effect of different cultivation conditions on the yield of yeast, taking into account different soil conditions. As a result of this interactive process, the biotechnology company/farmer is provided with a prediction machine (e.g. in the form of a computer system or expert system) that is trained to make predictions of the yield of yeast for different combinations of soil and cultivation conditions.

Example 3 - Field of Industrial Production/Manufacturing

[0171] This example concerns the case of a company manufacturing combustion engines or motors. Different compositions of fuel can be used to power these motors. However, the precise choice of composition will affect the functioning of the motor: some compositions may be cheap but cause the motor to deteriorate, cutting short its lifespan; certain other compositions may extend the lifespan of the engine but at the cost of increased energy consumption. The effect of the fuel composition on the engine's performance/lifespan may be particularly critical in the case of large engines used in factories and the like. It is desirable to know how variation of the fuel composition affects variables such as the lifespan of the engine, the fuel consumption, etc.

[0172] Clearly, tests involving determination of factors such as an engine's lifespan for different compositions of fuel can be extremely costly. Once again, it is desirable to be able to reduce the number of tests that must be performed in order to learn the effect of different fuel compositions on the engine and its performance.

[0173] The action-selection technique of the present invention can be used in order to determine the sequence of actions which should be performed so as to learn efficiently about the effect of different fuel compositions on the engine's performance. More particularly, below there is described the case where the engine manufacturer uses an action-selection system 1 as shown in Fig.4 to select actions in a sequence that will investigate the performance of his engine when supplied with different fuel compositions.

[0174] In this example, the percentage by weight of different components ***a, b,*** **...m,** in the fuel composition can be specified using parameters C_a, C_b, ... C_m, respectively. The amount of fuel consumed by the engine per hour of operation can be specified using a parameter ***FC***, the power output rating of the engine can be specified using a parameter ***P***, and the lifetime of the engine can be specified using a parameter, ***LT***.

[0175] In this example, the action-selection system 1 according to a preferred embodiment of the invention may be configured to recognise the parameters C_a, C_b, ... C_m as action variables. In this example, the action-selection system 1 will be configured such that its prediction module(s) 12 seek to predict what will be the useful lifespan of the engine ***(LT)***, the engine's fuel consumption (***FC***) and the power rating of the engine (***P***). In other words, in this example the set of predicted variables ***VAR*** includes three variables. This is another example in which no system-state/context parameters are defined for the action-selection system 1.

[0176] The action-selection system 1 according to the preferred embodiment of the invention will predict ***LT, FC*** and ***P*** for each of a number of candidate actions, that is, for each of a number of different possible fuel compositions.

[0177] Once again, on a majority of occasions the candidate action that the action selector 4 selects for performance, based on the output from the relevant knowledge-gain assessment modules 14, will be the one that is expected to produce the greatest increase in knowledge compared to the last occasion on which the system had a combination of values of C_a, C_b, ..., and C_m that are in the same region of system-state/context/action space as the combination specified in the candidate action. On a minority of occasions, a random action will be selected. Output is given as to which combination of values of C_a, C_b, ... C_m has been selected.

[0178] The engine-manufacturing company receives the output from the action-selection system 1, and may then operate an engine under the specified conditions and note the useful lifetime of the engine, the fuel consumption and the power output when so operated. The measured data ***LT, FC*** and ***P*** is input to the action-selection system as feedback (actual values of ***VAR***), enabling the calculator 5 of the action-selection system 1 to determine what was the actual error on the prediction made by the relevant prediction module 12 for this candidate action (this fuel composition).

[0179] As in the above-described first and second examples, as the series of tests progresses the action-selection system 1 will create new expert units (with respective prediction modules) which sub-divide system-state/context/action space according to a tree-like structure which reflects the relationships that exist between the different components of the fuel composition, the power rating of the engine, the fuel consumption of the engine and the lifetime of the engine.

IPR2025-00943
Tesla EX1002 Page 1500

[0180] At the end of the series of tests, the combination of the region identifier 2 and the prediction modules 12 of the action-selection system 1 has developed into a prediction machine that is trained to predict how the lifetime of the engine, its fuel consumption and power output are affected by the composition of the fuel supplied to it. When selling an engine, the manufacturer can supply this trained prediction machine to the end-user, enabling that user to evaluate the engine's performance for different possible fuel compositions.

[0181] It will be seen that, in this example of an application in the field of industrial production, the first aspect of the invention has provided the engine manufacturer with a machine which selects the actions in a sequence which enables the company to learn the effect of different fuel compositions on the performance of its engine.

[0182] Moreover, in this example of an application in the field of industrial production, the second aspect of the invention provides the engine manufacturer with a prediction-machine training system which, in an interactive process, trains a prediction machine so that it can predict the effect of different fuel compositions on the engine's performance. As a result of this interactive process, the engine manufacturer/engine user is provided with a prediction machine (e.g. a computer system or expert system) that is trained to make predictions of the engine's fuel consumption, power output and lifetime when using different fuel compositions.

Example 4 - Financial Field

[0183] Consider an example in which a person wishes to be able to predict what will be the value of a particular financial object (e.g. a building or other asset; shares, a warrant, an option or other financial instrument; etc.) thirty days after purchase.

[0184] The value of the financial object will be affected by prevailing economic conditions; these can be considered to be the system's context and can be described using context parameters such as the end-of-day level of the local stock market index (***SMI***), the current interest rate **(*IR*)**, etc.

[0185] The value of the financial object will often also be affected by a number of parameters defining that object. This can be understood by considering an example such as a (call) share option. This share option gives its owner the right to buy a certain quantity of shares on a particular day at a particular price, known as the "strike price". (An option to sell shares would be known as a "put" option). The share option itself can be traded and so will have a price; this is affected by the time left until the date on which the option can be exercised as well as the difference between the strike price and the actual market price for the same shares. On any given day it is usually possible to buy different options on the same shares, notably options having different strike prices and different exercise dates. The actual share price at the time of purchase of the option, and the price of the share option, can be considered to define the "state" of this kind of financial object and could be described using respective state parameters CP (current price of shares on date of purchase of option) and OP (current price of option).

[0186] In order to find out how the value of a call share option 30 days after purchase varies with economic conditions and the "state" of the option itself, it would be necessary to buy a number of options having different strike prices and exercise dates, in circumstances where the economic conditions (including the underlying share price) vary. Clearly each such purchase involves costs and, potentially, financial loss. Thus, it is desirable to be able to discover the relationships inherent in this system in an efficient manner. Use of the action-selection system according to the present invention enables this to be achieved.

[0187] Preferably, the user employs a fully-automated arrangement in which the action-selection system of the invention is adapted to be able to access online financial databases, electronic marketplaces and/or electronic trading systems (e.g. by connecting itself to websites), so as to be able to gather for itself the current values of the system-state/context parameter values, to initiate share option purchases and to obtain feedback on the outcome of those purchases.

[0188] An example of such a fully-automated system is shown, schematically, in Fig.7. In the case illustrated in Fig. 7, an action-selection system according to the first aspect of the invention is capable of accessing Internet financial databases (for example the NASDAQ website) in order to obtain current values for various economic indicators and for parameters defining potential products that could be purchased. The action-selection system is also capable of making purchases in the international financial marketplace (e.g. by making use of an electronic trading system or online trading service) and of obtaining feedback from that marketplace, automatically, with regard to the outcome of those purchases.

[0189] In this particular example, the action-selection system 1 shown in Fig.4 may be used in the fully-automated assembly of Fig. 7 and may be configured to recognise the end-of-day level of the local stock market index (***SMI***), the current interest rate (***IR***), the current share price (***PP***) and current option price (***OP***) as system-state/context parameters. The share option exercise date (***ED***) or time until exercise (***TE***), as well as the strike price (***SP***) are action variables.

[0190] The region identifier 2 and actual prediction error calculator 5 receive as inputs the values of the end-of-day level of the local stock market index (***SMI***), the current interest rate (***IR***), the current share price *(**PP**)* and current option price (***OP***) at a given time, t, when a purchase is going to be made. The prediction module(s) 12 of the action-selection system are configured to predict the value of the share option 30 days after purchase (based on the assumption that the values of the various system-state/context parameters that are input at time ***t*** will still hold when the actual purchase

IPR2025-00943
Tesla EX1002 Page 1501

is made). The latter assumption is reasonable, given that the action-selection system 1 can select an action, and output data identifying that action, within a few milliseconds, and economic variables such as share prices vary much more slowly than that.

[0191] In this example the action-selection system 1 of the invention may be configured to process action variables ED and SP, corresponding to tuning of the exercise date and strike price parameters.

[0192] Once the action-selection system 1 has been configured, it accesses financial databases or the like so as to input details of the current values of the stock market index, ***SMI***, the official central bank interest rate, ***IR***, and the current share price, ***PP***. The action-selection system 1 may itself include a candidate action generator so as to generate possible combinations of values for the exercise date and strike price, ***ED*** and ***SP***. Alternatively, the different sets of candidate actions may be generated by a separate module or by the user. The action-selection system 1 accesses finanacial databases or the like in order to determine the prices, ***OP***, of call share options having the combinations of exercise dates and strike prices specified in the candidate actions.

[0193] For each of these different potential call-share-option purchases, the action-selection system 1 of the invention seeks to predict what will be the value of the bought option in 30 days time, taking into account the input values for the economic indicators (***SMI*** and ***IR***), the current share price (***PP***), the current price of the option (***OP***), and the time until exercise of the option (***TE***) which the system can calculate for itself, provided that it knows the current date.

[0194] As an alternative to using a fully-automated arrangement, the user could make use of the action-selection system 1 only to select actions, and himself input the values of the various system-state/context parameters, make the selected purchases, and provide feedback as to the outcomes.

[0195] The action selector 4 of the action-selection system will select which action (purchase) should be tried next. In general, the selected action will be the one which is expected to produce the greatest increase in knowledge (assessed "locally" in system-state/context/action space), although random actions will be selected with probability ***p***. The share option having the selected combination of strike price and exercise date can then be bought automatically for example via an electronic trading system (or by the user). The value of the purchased share option after 30 days is noted and fed back to the action-selection system automatically, from internet financial databases for example, or manually by the user.

[0196] Initially, the action-selection system 1 of the invention only recognises one region in context/system-state/action space and so all predictions are made a common prediction module. However, as the purchases suggested by the system are made, and feedback is received on the actual value of the purchased share options 30 days after purchase, the system will begin to divide up the space into different regions, creating new expert units (and corresponding new prediction modules 12) as required.

[0197] After a certain number of purchases have been made, provided that there is some non-random element in the system, the action-selection mechanism 1 of the invention will have developed expert units having prediction modules 12 that are tailored to the different situations that can arise (different combinations of economic indicators, share price, option parameters, etc.). The series of actions (purchases) can then be stopped and the user has available to him a machine which can predict what will be the value in thirty days time of share options defined by different combinations of strike price and exercise date, taking into account the current economic context and share price. He can compare these predicted values with the respective current prices of the share options in order to decide which purchases are more likely to provide financial gain.

[0198] The trained action-selection system 1 could be sold to third parties, as a trained option-value prediction machine. If the third party acts on the suggestions made to it by the trained prediction machine, he can feedback the results of his purchases (i.e. actual option value after 30 days) so as to provide the prediction machine with further examples to process. In other words, even after the prediction machine has been trained it is still possible to generate further "training examples" and refine the operation of the prediction machine yet further.

[0199] Incidentally, in this particular application it is not essential for the automated system (or the user) to actually make the selected purchases (perform the selected actions). The price of the selected share-option can be determined 30 days after it would have been purchased, and this price fed back to the action-selection system 1. It is true that the actual purchase of the share-option could have had an effect on the market value quoted for options having the same specification. However, unless an extremely large purchase was intended (or the market in shares of this type is extremely small) the absence of an actual purchase will no have a significant effect on the market value for this option in 30 days time.

[0200] It will be seen that, in this example of an application in the financial field, the first aspect of the invention provides a user with a machine which selects the actions in a sequence which enables the user to learn the effect of making different purchases in different conditions. The first aspect of the invention also provides a fully-automated arrangement in which the action-selection system 1 obtains for itself the system-state/context parameter values it requires (for example, by connecting itself to on-line financial databases), determines which actions should be performed, optionally performs them (by connecting itself to electronic trading systems), and obtains feedback on the outcomes (for example, by once again connecting itself to online financial databases). In this fully-automated arrangement the action-selection system 1 could be considered to comprise a software robot (or "softbot") making the required connections to websites/trading

IPR2025-00943
Tesla EX1002 Page 1502

systems.

[0201] Moreover, in this example of an application in the financial field, the second aspect of the invention provides the user with a prediction-machine training system which, in an interactive process, trains a prediction machine so that it can predict the effect of making different purchase in different conditions. As a result of this interactive process, the user is provided with a prediction machine (e.g. in the form of a computer system or expert system) that is trained to make predictions of the future value of a particular financial instrument given different combinations of conditions at the time of purchase.

THE THIRD ASPECT OF THE INVENTION

[0202] According to the third aspect thereof, the present invention provides a technique and system for managing learning by, or organizing behaviour of, a self-developing device or system, as well as a self-developing device or system making use of such a technique.

[0203] Some types of self-developing devices are robotic devices, or other devices which have actuators and sensors enabling them to interact with their environment. According to the third aspect of the invention, it is possible to manage learning by such a self-developing device, and/or organize its behaviour, by treating that self-developing device interacting with its environment as a system ***S***, and investigating its behaviour in a series of actions that are chosen in an autonomous fashion according to the action-selection technique of the first aspect of the invention. In general, the system-state/context parameters of such a system are the signals received by the self-developing device's sensors and the action variables will be the setting of the self-developing device's actuators. In this case, the system-state/context/action space equates to the self-developing device's sensory-motor state-space.

[0204] However, it is to be understood that the third aspect of the present invention is not limited to application with single self-developing devices which have sensors and actuators.

[0205] More particularly, the present invention is also applicable to systems which contain distinct separate components or entities working together, in which case the selected action may be the setting of respective values for action variables of one or more of the system components/entities.

[0206] Furthermore, the third aspect of the present invention is applicable to self-developing devices or systems which can be considered to be passive, that is, which do not have physical actuators and may not even have sensors. For example, the present invention can be applied to a self-developing computer system which receives or generates certain inputs, and produces outputs, thereby interacting with an environment. According to the third aspect of the invention, it is possible to manage the learning of such a self-developing system, or organize its behaviour, by treating that self-developing computer interacting with its environment as a setup or system S, and investigating the behaviour of this setup in a series of actions that are chosen in an autonomous fashion according to the action-selection technique of the first aspect of the invention. The system-state/context parameters in this case will generally be the inputs received or generated by the self-developing computer, and the set of action variables will be one or more variables whose value (s) are set based on one or more of the outputs from the self-developing computer.

[0207] For example, for a self-developing system constituted by a computer, the system-state/context parameters could include: data indicating the state of certain internal elements of the computer (e.g. status of flags, full/empty status of buffers, existence of a state of overflow, etc.), data input by a human operator, data indicating the state of a context or environment with which the computer system interacts, etc. This latter environment could be physical, for example in a case where the computer system controls a physical device or system, or it could be non-physical, for example in a case where the computer system is trying to teach itself what patterns there are in the data it is processing. In either case, management of learning by the computer system, or organization of its behaviour, can be achieved using the method according to the third aspect of the present invention.

[0208] Moreover, the third aspect of the present invention can be used to manage learning by, or organize the behaviour of, an expert system. For example, an expert system may be intended to interact with human users, by asking them a series of questions in order to make some determination (for example, as to the identity of the user). However, before using the expert system for that purpose, it may be useful for the expert system to perform a series of tests, in which it tries out different possible questions or sequences of questions in order to determine which sequence is most likely to enable it to make accurate determinations later on. Such a series of tests could be selected using a technique according to the third aspect of the invention, in order to organize the behaviour of the expert system/manage its learning during the test phase. In this example application, the action variables would be different possible questions that could be posed.

[0209] The third aspect of the invention will be described below with reference to a simple experimental set-up. As shown in Fig. 8, this experimental set-up involves a simple robot, RBT, moving in a square space that is enclosed by walls, WL. The robot is formed of a box mounted on a left wheel and a right wheel, and it can control the speed of each wheel separately. The robot also has two infrared sensors at the front, enabling it to sense the walls, a "toy sensor" to be discussed in greater detail below, and can produce sounds at a selected one of three frequencies (labelled ***f1***, ***f2*** or ***f3***). It can be considered that the robot has two wall sensors (the two infrared sensors), one toy sensor, and three

IPR2025-00943
Tesla EX1002 Page 1503

actuators (the controller of left wheel speed, the controller of right wheel speed, and the controller of sound frequency). The robot has an internal clock and updates the readings from its sensors, as well as the settings of its actuators, once every 1000 ms.

[0210] In the example discussed here, the robot has a decision-making unit integrated into it which determines which action the robot should take at a given time step. This decision-making unit comprises an action-selection system according to the first aspect of the invention, for example the action-selection system 1 of Fig.4. In the present case, this action-selection system can be considered to be a module that manages the robot's learning process thereby obtaining a particular kind of organization of the robot's behaviour. It should be noted that the invention is not limited to the case where the action-selection system according to the first aspect of the invention (or the decision-making unit) is on-board the robot. This system and unit could be separate from the robot but in communication with it.

[0211] Furthermore, for simplicity it is assumed in the following discussion that the decision-making unit bases its decisions on which action the robot should take solely on the selections made by the action-selection system according to the first aspect of the invention - in other words, in this example there is only one drive motivating the robot's actions.

[0212] In the experimental set-up of Fig.8 there is an intelligent toy, T, in the same space as the robot. The toy moves according to the sounds produced by the robot. In particular, when the robot emits sounds at frequency ***f1***, the toy moves randomly around the space at a speed higher than the robot can move; consequently it seems to the robot as if the toy is appearing and disappearing at random positions. When the robot emits sounds at frequency ***f2***, the toy stops moving. When the robot emits sounds at frequency ***f3***, the toy jumps into the robot. The robot can detect the toy, here using a vision system, and calculate the toy's position relative to the robot, in terms of x and y co-ordinates. This vision-based system for detecting the toy and determining its position is designated the "toy sensor".

[0213] The above-described experimental set-up was simulated using the well-known "Webot" mobile robot simulation software provided by Cyberbotics Sarl of Lausanne, Switzerland (available at www.cyberbotics.com).

[0214] At each time step, the robot (or, strictly speaking the action-selection system thereof) is arranged to decide what action to take: that is, what settings should be given to the speed of the left wheel, the speed of the right wheel, and the frequency of the sound the robot produces. Thus, in this case the action variables of the system are the left-wheel speed, right-wheel speed and sound frequency produced by the robot.

[0215] Simulations were performed in which the robot attempted to predict what would be the distance to the toy one time step (i.e. 1000 ms) after the robot took an action in a particular context (i.e. for a particular set of sensor readings). After elapsing of the 1000 ms period, the robot determined the actual position of the toy relative to the robot's own position in terms of (x,y) co-ordinates, using the toy sensor. Based on the sensed co-ordinate values, the robot calculated the actual distance to the toy.

[0216] When the robot produces the sound at frequency ***f1***, the toy moves around randomly. In this case, it is impossible for the robot to determine a relationship which will accurately predict the distance between the robot and the toy. There is no correlation between the distance to the toy and the robot's sensor values and wheel speeds. Nothing worthwhile can be learned by exploring this situation. In other words, the robot will not learn efficiently if it spends a long time emitting sound at frequency f1.

[0217] When the robot produces the sound at frequency f3 the toy jumps inside the robot and so the distance between the robot and the toy is zero. It is fairly easy for the robot's action-selection system to learn this relationship between production of sound at frequency f3 and zero distance to the toy (even in the presence of noise, for example due to the sensors). This relationship has very low complexity.

[0218] When the robot produces the sound at frequency f2 the toy stops moving, but because of the robot's own movements it will still seem to the robot as if the toy *is* moving. Nevertheless, if the robot makes a series of observations while continuing to emit sound at frequency f2, its action-selection system should eventually learn that the distance to the toy is related to the robot's own movements.

[0219] The interval between selecting an action and observing its outcome was set at 1000 ms, which is a relatively long period of time. Accordingly there could be substantial changes in the toy's position relative to the robot. This made it more difficult for the robot's action-selection system to learn the above-mentioned relationships that exist between the motor speeds of the robot's wheels, the sound frequency produced by the robot and the distance to the toy.

[0220] More particularly, in this example the goal was for the robot's action-selection system to learn mappings to enable it to predict, accurately, at a time ***t*** what would be the distance to the toy in 1000 ms time (i.e. at ***t***+1) IF the robot takes a particular action A(t) in the robot's current sensory-motor context ***SC(t)***. The values of the system-state/context parameters and action variables in this case are the set of values measured by the robot's sensors (wall sensors and toy sensors) at the time ***t***, as well as the settings of the robot's actuators at the time ***t***.

[0221] The only real complexity in the mapping that the robot's action-selection system must learn resides in that part of the mapping which relates to production of sound at frequency f2. When sound at frequency f2 is produced, the distance to the toy varies, as a function of the sensory context SC(t) and the actions that are performed **A**(***t***), in a complicated but learnable manner.

[0222] If the robot is to learn efficiently it should spend a lot of time exploring situations in which sound is produced

IPR2025-00943
Tesla EX1002 Page 1504

at frequency f2 (where there is a lot to learn), a little time in situations where sound is produced at frequency f3 (where there is a small amount to learn), and practically no time in situations where sound is produced at frequency f1 (where there is nothing that can be learned). The third aspect of the present invention provides a technique which enables the robot to autonomously organize its own behaviour so as to focus attention appropriately on the situations which will render its learning efficient.

[0223] Fig.9 shows, schematically, the main components of a preferred embodiment of action-selection system 21 that can be used to manage learning by, and/or organization of the behaviour of, the robot of Fig.8. It will be noted that this action-selection system 21 has a similar construction and functioning to that of the action-selection system 1 that is illustrated schematically in Fig.4. Accordingly, the same reference numerals are used in Fig.9 to designate elements having a similar function to those in the action-selection system 1 of Fig.4.

[0224] Once again, it is to be understood that this action-selection system 21 would probably be implemented in software and the various "modules" shown in Fig.9 are provided for illustrative purposes only. The various functions performed in the action-selection system 21 may be split between modules in different ways from that shown in Fig.9.

[0225] The action-selection system 21 of Fig.9 receives inputs from sensors and actuators, these inputs representing the robot's sensory-motor state. In the present example, the inputs will be from the robot's infrared wall sensors, from the toy sensor, and from the robot's actuators (signals indicating the values of right wheel speed, left wheel speed and sound frequency - these signals may be supplied from controllers controlling the wheels/sound generator, rather than from the wheels/sound generator themselves).

[0226] In the preferred embodiments according to the third aspect of the invention, the signals from each of the sensors and actuators are normalised so as to fall within an interval 0.0 to 1.0. In the case of the sound frequency produced by the robot, f1 is produced when the "frequency" variable is in the range 0 to 0.33, f2 is produced when the "frequency" variable is in the range 0.34 to 0.66, f3 is produced when the "frequency" variable is in the range 0.67 to 1.0.

[0227] The action-selection system 21 includes a candidate action generator 25 which generates a list of possible actions that the robot could perform next, given the current values of its system-state/context parameters $\boldsymbol{SC}(t)$.

[0228] For each candidate action, the prediction module 12 in the relevant expert unit of the knowledge-gain assessment module 3 makes a prediction, $\boldsymbol{d}(t+1)$, of the distance to the toy at the next time step (in 1000 ms time). For each prediction made by a prediction module 12, the meta prediction module 13 of the same expert unit produces an estimate, $\boldsymbol{Ep}(t+1)$, of what will be the error in the predicted distance to the toy. The knowledge-gain assessment module 14 of this same expert unit compares the expected prediction error for the candidate action with actual prediction error data that has been observed and stored in the memory 15 of this expert unit, and outputs to the candidate action selector 4 a signal indicating the knowledge gain expected to arise if this candidate action is performed.

[0229] The action selection module 4 receives the expected-knowledge-gain signals, selects which action should be performed next and outputs appropriate control signals to the robot's actuators. The action selection module 16 also indicates to the calculator 5 which candidate action has been selected.

[0230] The calculator 5 calculates what is the actual prediction error, $\boldsymbol{E}_{\boldsymbol{Ract}}(t+1)$, and stores this data in the memory 15 of the appropriate expert unit.

Experimental results

[0231] A number of simulations were performed to demonstrate the efficiency of the technique of the third aspect of the present invention for managing learning in a self-developing device. These simulations used the experimental set-up of Fig.8 and ran for a large number of time steps. In these simulations the robot's action-selection system had no prior "knowledge", in particular it was not programmed to know that there is a difference between setting the speed of the wheels and setting the sound frequency - these parameters were, in effect, variables of unidentified type as far as the robot was concerned. The robot's action-selection system was designed to treat as an action variable any variable that it could tune, and to treat as a system-state/context parameter any parameter for which it could determine a value.

[0232] Fig. 10 is a graph indicating how actual prediction error varied over time (as the number of performed actions increased) when the action-selection system 21 of the preferred embodiment of the invention was used to help the robot of Fig.8 to learn to predict the distance to the toy T. In this simulation, the action selection module 4 had a 0.85 probability of selecting the candidate action expected to produce the greatest knowledge gain and a 0.15 probability of selecting a random action.

[0233] In Fig.10, a separate line is drawn to represent the evolution of actual prediction error for each "leaf" expert unit. On splitting of an expert unit, each of the new "leaf" expert units inherits from the initial expert unit those training examples (and actual prediction error values) which relate to the sub-region of system-state/context/action space now assigned to this new "leaf" expert. Thus, the prediction error traces in Fig.10 overlap at the left-hand side of the graph and progressively split from each other at the times when the corresponding expert units split. Thus, a tree structure appears in Fig. 10 and this corresponds to the decision-tree structure developed by the region identifier 2.

[0234] It will be seen that three groups of expert units developed during the simulation. Tests confirmed that the

IPR2025-00943
Tesla EX1002 Page 1505

different branches of the decision-tree structure did indeed correspond to different situations; the topmost branch in Fig. 10 corresponds to situations in which the robot produces sounds at f1, the middle branch corresponds to situations in which the robot produces sounds at f2, and the lowest branch corresponds to situations in which the robot produces sounds at f3.

[0235] As shown in Fig. 10, the decision-tree structure quickly developed so as to distinguish those situations in which sound was being produced at f3. After only 250 time steps the robot (or, rather, its action-selection system 21) was able to differentiate situations in which sound was being produced at f3 from situations where sounds were being produced at other frequencies. At this time, the initial (or "root") expert unit in the action-selection system 21 was split into a first "branch" expert unit handling situations where the "sound frequency" variable took values 0.67 to 1.0 (i.e. sounds at f3) and a second "branch" expert unit handling situations where the sound frequency variable took values 0.0 to 0.66 (i.e. sounds produced at f1 or f2).

[0236] The robot discovered that, initially, the situations with sounds at f3 were the most important sources of learning progress. Thus, the majority of the time actions were selected which kept the robot in situations where the robot produced sounds at frequency f3. However, the robot very quickly learned all there is to know about such situations, namely that the distance to the toy is substantially zero in these cases. Accordingly, the prediction error quickly fell to around zero and these situations were no longer a source of learning progress.

[0237] During the initial stage of the simulation, although the robot preferred to explore situations with sound at f3 (because these offered the greatest learning progress) it nevertheless performed some actions which involved production of sounds at f1 or f2 because of the random selection of actions by the action-selection module 4. After a time, the "branch" expert unit handling situations where the sound frequency variable took values 0.0 to 0.66 (i.e. sounds produced at f1 or f2) had accumulated sufficient examples to meet above-mentioned criterion C1. This occurred at around the 320th time step, at which time this branch expert unit split into an "upper branch" expert unit handling situations where the sound frequency variable took values 0.0 to 0.33 (i.e. sounds produced at f1) and a "central branch" expert unit handling situations where the sound frequency variable took values 0.34 to 0.66 (i.e. sounds produced at f2).

[0238] At about the 400th time step the robot (or, strictly speaking, the knowledge-gain assessment modules of the expert units in the action-selection system 21) accurately determined that situations with sounds at f2 would provide greater learning progress than situations with sounds produced at f3 or situations with sounds produced at f1. Accordingly the action-selection module 4 tended to select actions which would involve production of sounds at f2.

[0239] Because the action-selection module 4 occasionally chooses random actions, the robot sometimes performs actions involving production of sounds at f1. However, when sounds are produced at f3 the expert units cannot learn anything to help them improve the accuracy of their predictions of the distance to the toy. Accordingly, the expert units quickly learn that actions resulting in the production of sound at f3 can produce no knowledge gain. It will be seen in Fig. 10 that the branch corresponding to situations with sounds at f1 ceases to be explored.

[0240] On the other hand, as more and more training examples are accumulated for situations in which sounds are produced at f2, the expert units can begin to see relationships between the robot's actions and the distance to the toy, thus enabling their predictions to improve. Accordingly, the robot explores different situations in which sound is produced at f2, experimenting with varying the speed of its left and right wheels. It has been found that, when using the active-learning management algorithm of the preferred embodiment of the present invention, this exploration of motor speed settings is organized and the speed settings are themselves selected in a manner which optimizes learning progress.

[0241] For the above-described simulation, it is interesting to consider how much time the robot spends in situations where sounds are produced at f1, f2 and f3, respectively. In order to investigate this, the simulation was repeated and Fig. 11 indicates how the percentage of time the robot spent in performing actions of different kinds varied as time went on (after the robot had performed a greater and greater number of actions). In Fig.11, separate traces represent the percentage of time spent performing actions in which sound is produced at f1, f2 and f3, respectively.

[0242] Figure 11 shows that the robot initially concentrates its attention on situations in which sounds are produced at f3. After the robot has mastered these situations it turns its attention to situations in which sounds are produced at f2.

[0243] From a consideration of the results of these simulations it will be seen that, by managing the robot's learning using the preferred embodiment of the third aspect of the invention, the robot has been able, autonomously, to scale the complexity of the situations it investigates, starting with situations of relatively low complexity and moving on to those of greater complexity. It has also been able to avoid spending time exploring situations where there is nothing to be learned. Thus, the third aspect of the present invention enables the robot to engage in efficient learning. Moreover, use of this technique also enables the robot to organize its own behaviour without outside intervention. This is a highly-desirable achievement.

[0244] As mentioned above, the active-learning management technique of the third aspect of the invention can be termed an Intelligent Adaptive Curiosity drive.

IPR2025-00943
Tesla EX1002 Page 1506

Comparative Examples

[0245] Figs. 12 to 17 are graphs representing the results of simulations that were performed for the purposes of comparing the intelligent adaptive curiosity drive of the third aspect of the present invention with other known techniques for managing active learning.

First Comparative Example

[0246] Fig.12 shows how prediction error evolved over time in a simulation comparable to that describe above, in which the robot used expert units handling different regions of sensory-motor state-space and progressively divided up that space into regions according to the same criteria C1 and C2. However, in the case of this first comparative example, there was no assessment of expected knowledge gain; instead, actions were selected entirely at random. The algorithm used in this first comparative example can be denoted the "random" algorithm (RAN).

[0247] Fig. 12 shows a typical example of results obtained for the RAN algorithm after 15000 random actions. Sensory-motor state-space has been divided into three main regions, respectively corresponding to situations in which sounds are produced at f1, f2 and f3. Early on, situations have been explored in which sounds are produced at f3 and the system has learned how to predict the distance to the toy in those situations. Moreover, over time the prediction error associated with situations in which sound is produced at f2 has decreased. However, the system continues to perform actions which involve the robot producing sounds at f3. It is completely unproductive to do so, because the robot can learn nothing in these situations. Thus, in contrast to the intelligent adaptive curiosity drive of the present invention, the RAN algorithm does not produce efficient learning.

Second Comparative Example

[0248] Fig. 13 shows how prediction error evolved over time in another simulation comparable to the first described above. Once again the robot used expert units handling different regions of sensory-motor state-space and progressively divided up that space into regions according to the criteria C1 and C2. However, in the case of this second comparative example, actions were selected according to a conventional approach that can be summarized "choose the action which you expect to maximize prediction error". This can be termed the "maximizing" algorithm (MAX). It will be seen that this algorithm seeks out novel situations (on the assumption that these situations offer the greatest potential for learning).

[0249] Fig. 13 shows a typical example of results obtained using the MAX algorithm. Once again, sensory-motor state-space has been divided into three main regions corresponding to situations with sounds produced at f1, f2 and f3, respectively and, at an early stage, the system learns how to predict the distance to the toy in situations where sounds are produced at f3. However, situations in which sounds are produced at f2 and f3 are not explored in detail; instead, the system concentrates on situations where sounds are produced at f1. This is made even clearer by considering Fig. 14, which shows the percentage of time the robot spent in the three different types of situations. After only 1000 time steps, the robot is already spending 35% of its time in situations in which sounds are produced at f1.

[0250] It can be seen that, when seeking to manage active development, the MAX algorithm is counter-productive: it leads the robot to explore situations in which nothing can be learned.

Third Comparative Example

[0251] Fig.15 shows how prediction error evolved over time in a fourth simulation. This third comparative example was similar to the second comparative example except that actions were selected according to a conventional approach that can be summarized "choose the action which you expect to minimize prediction error" (a "minimizing" algorithm (MIN)).

[0252] In fact, Figs. 15 and 16 are the same sort of graphs as Figs.10 and 11 but, this time, plotted for a simulation using the MIN algorithm. In this case the algorithm drives the robot to stay in situations in which sound is produced at f3, the most easily learnable and predictable situations. This is not as inefficient as the MAX algorithm, but it is still counter-productive, very quickly the robot will have learnt all there is to know about the simple situations but it will not progress to explore more complex situations.

Advantages of the Third Aspect of the Invention

[0253] The action-selection technique according to the third aspect of the present invention allows a robot (or other self-developing device or system) to scale the complexity of its learning situations by successively and actively focusing its activity on problems of progressively increasing difficulty. This is achieved through the efficient and operational discovery, within system-state/context/action space, of regions of varying complexity, then evaluation of the relative

IPR2025-00943
Tesla EX1002 Page 1507

complexity of these regions (in terms of learning potential), and the exploitation of this information by the robot (or other self-developing device or system) to self-organize its behaviour.

[0254] Other techniques for managing learning do not lead to development of this kind of organized behaviour except, perhaps, in very particular kinds of environment: for example, in an environment in which perturbations of increasing complexity are continuously generated. Such environments can be produced by human intervention but, otherwise, they are believed to be extremely rare. By way of contrast, as long as there is some potential complexity in the robot/environment, the third aspect of the present invention can efficiently manage open-ended learning without complicated dynamics of this kind.

[0255] The efficiency of the active learning achievable using the third aspect of the present invention is greater than that obtained using the algorithms of the comparative examples. This was shown by a further set of simulations which involved computing, at different stages in the simulation, how much the robot has learned about the more interesting situations, that is, the situations in which sounds are produced at f2.

[0256] In this set of simulations, the robot used the active-learning management technique of the preferred embodiment of the invention, the RAN, MIN and MAX algorithms, respectively, to predict the distance to the toy at each of a large number of time steps. After every 150 time steps, each simulation was temporarily stopped and the robot was exposed to a test database to see how well the robot could predict the distance to the toy for a number of situations in which sounds were being produced at f2. When each simulation was restarted, the examples in the test database were not maintained in memory (so as to avoid biasing the test results).

[0257] Incidentally, the test database was generated as follows. The robot was operated such that it selected actions at random. For each situation in which sounds were produced at f2, a note was made of the robot's sensory-motor data at that time as well as the distance to the toy.

[0258] Fig. 17 shows the results of this set of simulations, more especially it shows, for each of the different algorithms, how the mean prediction errors produced on analysing the test database varied over time.

[0259] It will be seen from Fig.17 that, initially, the RAN algorithm gives the lowest prediction error. In other words, early on it is the random algorithm that lets the robot learn the most about the situations where sounds are produced at f2. However, after about 3000 actions (time steps), the active-learning management technique according to the third aspect of the present invention begins to produce markedly lower prediction errors for these situations with sound at f2. This corresponds to the fact that the robot has turned its attention away from situations having sound at f3 (which the robot has thoroughly explored) and towards the situations with sounds at f2. In fact, the robot is spending 85% of its time exploring situations with sounds at f2 and rapidly learns how to improve its predictions of the distance to the toy in such situations.

[0260] After about 3000 time steps (actions) the active-learning management technique according to the third aspect of the present invention is the algorithm which has enabled the robot to learn most about the situations in which sounds are produced at f2. The mean prediction error using this algorithm stabilises at a low value of around 0.090. The mean prediction error produced by the MAX and MIN algorithms, respectively, remains at a substantially higher level (around 0.11). The mean prediction error produced by the RAN algorithm does, eventually, decrease to stabilise at a low value. However, this low value is around 0.092 which is still higher than the mean prediction error produced by the algorithm according to the present invention. Moreover, this low value is attained by the RAN algorithm after around 11000 actions. In other words, both the RAN algorithm and the algorithm according to the present invention enable the robot to learn about situations in which sounds are produced at frequency f2, but this learning is about 2.6 times faster when using the algorithm according to the third aspect of the present invention.

[0261] Although the various aspects of the present invention have been described above in terms of embodiments which are presently preferred, it is to be understood that the invention is not limited by the particular details of the preferred embodiments. In particular, the skilled person will readily appreciate that various changes and developments can be made in the preferred embodiment, within the scope of the annexed claims.

[0262] For example, although the prediction modules in the preferred embodiments of the invention use a nearest-neighbour prediction algorithm, the invention is not limited to the case where this algorithm is used. Similarly, although the action-selection module 4 of the preferred embodiments of the invention chooses random actions with a probability of 0.15, the invention is not limited to use of this probability value.

[0263] Also, although the region identifier used in preferred embodiments of the invention splits a region ***R*** into two or more smaller regions when the number of training examples in region ***R*** has reached a value *NS*=250, this number is merely illustrative, not limitative. Moreover, as indicated above, other splitting criteria could be used which do not count the number of training examples in a region.

[0264] Moreover, as indicated above, the distribution of functions between modules in the systems of Figs. 4, 5 and 9 is purely illustrative. Other distributions of the described functions between the modules are possible and/or a greater or lesser number of modules could be used to perform the same functions - for example, one module could be used for defining regions in system-state/action/context space and another module could be used for identifying which region applies to a particular combination of system-state/context parameters and action variables. Moreover, it is not necessary

IPR2025-00943
Tesla EX1002 Page 1508

to include distinct modules in the action-selection systems and training systems of the invention. The various functions of the preferred embodiments of the invention could be performed, for example, in a computer program, by one or more software modules having overlapping functions.

[0265] Also, the skilled person will readily appreciate that certain of the functions ascribed to the different modules in Figs.4, 5 and 9 can be replaced by other functions, or combinations of functions, having an equivalent result. For example, as shown in Fig.4, the action selector 4 sends the calculator 5 of actual prediction error a signal indicating which action was actually selected for performance. The calculator 5 then retrieves the prediction relating to the action that has been selected for performance. However, the system 1 could be arranged such that calculator 5 receives all predictions made by the prediction modules 12 for a given set of candidate actions and identifies which of these predictions to use when assessing actual prediction error by studying the signal received from the action selector.

[0266] Furthermore, although the above description refers to distinct expert units, each assigned to calculate the expected knowledge-gain for different regions of system-state/context/action space, and refers to old expert units being split, or to new expert units being "created", this is not limitative. More particularly, it is to be understood that the important factor is that the knowledge gain should be assessed in an appropriate manner for each respective region in system-state/context/action space, whether this is achieved using an expandable set of distinct expert units, or using a single device (or software module) configured to load the prediction function, meta-prediction function and stored training example data applicable to the relevant region in system-state/context/action space, or using some other means.

[0267] Moreover, systems embodying the present invention could include a specific module for collating and storing the training example data and supplying it, as needed, to the expert unit handling the appropriate region in system-state/context/action space.

[0268] Furthermore, as indicated above the intelligent adaptive curiosity drive provided by the present invention can be just one of a plurality of drives motivating the selection of actions by a self-developing device. In such a case, the action-selection system 21 of Fig.9 would not itself make the final selection of which action the robot would take. More especially, the action selection module 4 of Fig.9 could be omitted and the knowledge-gain values produced by the various knowledge-gain assessment modules 14 could be output to a central controller as reward values indicating how desirable the associated candidate actions are considered to be. The central controller would process these reward values together with others generated by other drives motivating the robot's behaviour in order to decide which action the robot should take next. The central controller would send a signal to the module 21 indicating which candidate action has been selected, thus enabling actual prediction error to be calculated and stored in the memory 15 of the appropriate expert unit ***EXP***.

[0269] Finally, it is to be understood that robotic devices, artificial intelligence systems, expert systems and the like are not limited to using only the third aspect of the invention. The first aspect of the present invention can be applied for selecting the actions to be taken by such devices and systems, even when these devices/systems do not form part of the setup S whose behaviour is being investigated. The second aspect of the present invention can be applied in order to train prediction machines included in these devices and systems, even when these devices/systems do not form part of the setup S whose behaviour is being learned.

Claims

1. An automated action-selection system (1) adapted to generate signals specifying values for a set of one or more action variables defining an action that can be taken whereby to affect a setup (***S***), the automated action-selection system (1) comprising:

   input means for receiving signals indicative of the value, at a time ***t***, of a set of system-state/context parameters (***SC(t)***) describing the state and/or context of the setup (***S***);
   a region definer (2) adapted to define a set of regions in a multi-dimensional system-state/context/action space, each dimension of the system-state/context/action space being defined by a respective different parameter or variable of the sets of system-state/context parameters and action variables;
   means for determining a set of candidate actions, each candidate action consisting of a possible set of values for the action variables;
   a region identifier (2) for identifying the region in system-state/context/action space containing the combination of a given candidate action with values of any system-state/context parameters at time ***t***;
   a prediction unit (12) adapted to predict the value of a set of one or more predicted variables (***VAR***) a predetermined interval after time ***t***, wherein a prediction function applied by the prediction unit (12) depends upon the region in system-state/context/action space containing the combination of this given candidate action with any system-state/context parameters at time ***t***;
   calculator means adapted to calculate, for selected candidate actions, a respective indicator of the expected

IPR2025-00943
Tesla EX1002 Page 1509

error in the prediction made by the prediction unit (12) for said selected candidate action;
memory means for storing indicators of actual prediction errors made by the prediction unit (12) for respective candidate actions selected on one or more previous occasions;
assessment means (14) adapted to evaluate the expected improvement in the performance of the prediction unit (12) if a given candidate action is performed, wherein an assessment performed by the assessment means (14) depends upon the region (***R***) in system-state/context/action space containing the combination of this given candidate action with the values, at time ***t***, of any system-state/context parameters, and the assessment means is further adapted to evaluate said expected improvement by comparing an indicator of said expected error in prediction with an indicator of the actual prediction error that existed on one or more occasions, previous to time ***t***, when the setup (***S***) had a combination of system-state/context parameters and action variables located in the same region (***R***) of the system-state/context/action space;
means for generating a signal indicating the desirability of selecting a given candidate action for performance, said signal being dependent on the expected improvement in the performance of the prediction unit (12) evaluated by the assessment unit (14) for said given candidate action; and
an action selector (4) adapted to select one of the candidate actions for performance based on said signals indicating the desirability of selecting a given candidate action and further adapted to output a signal indicating the selected candidate action.

2. The automated action-selection system according to claim 1, and comprising an action selector (4) adapted to select an action for performance to affect the setup (***S***), the action selector (4) is adapted to have a probability of ***p***, where $0 < p < 1$, of selecting that one of the set of candidate actions that the assessment unit (14) evaluates to be expected to yield the greatest improvement in performance of the prediction unit (12); wherein the action-selection system (1) is adapted to output data defining the action selected by the action selector (4).

3. The automated action-selection system according to claim 1 or 2, wherein the region definer (2) is adapted to divide an existing region (***R***) into two or more new regions when a first criterion (***C1***) is met.

4. The automated action-selection system according to claim 3, and comprising counting means for counting the number of occasions on which a first and second condition are both met, the first condition being that an action is taken affecting the setup (***S***), and the second condition being that the combination of the action variable values defining said action and the values of system-state/context parameter values at said time said action is taken falls within the region (***R***);
wherein the first criterion (***C1***) is met when the counting means has counted up to a predetermined number (***NS***).

5. The automated action-selection system according to any one of claims 1 to 4, and comprising meta prediction means (13) adapted to evaluate the expected error in the values of said set of variables (***VAR***) predicted by the prediction means (12), a prediction function applied by the meta prediction means (13) depending upon the region in system-state/context/action space containing the combination of said given candidate action with the system-state/context parameters at time ***t***; wherein the assessment means (14) is adapted to evaluate the expected improvement in the performance of the prediction means (12) if a given candidate action is performed by evaluating the decrease in prediction error that is expected to result from performance of said given candidate action.

6. The automated action-selection system according to any of claims 1 to 5, and comprising means for receiving feedback regarding the actual values of the predicted variables (***VAR***) resulting from performance of a given action; the prediction means (12) is responsive to the feedback data whereby to adapt the prediction function that is applied for candidate actions in the same region of system-state/context/action space as the region containing the combination of the action which produced the feedback data and values of any system-state/context parameters at the time when said given action was performed.

7. The automated action-selection system according to claim 6, and comprising memory means for storing training example data; for a given performed action said training example data comprising action variable values, any system-state/context parameter values applicable at the time said given action was performed, and feedback data defining the actual values of the set of predicted variables (***VAR***).

8. The automated action-selection system according to claim 7, wherein the region definer (2) is adapted to divide a region (***R***) of system-state/context/action space into two or more new regions, when a first criterion (***C1***) is met, by defining boundaries of the new regions so that there is a balance between the numbers of training examples in each new region, and the variance of the training examples is minimized in system-state/context/action space, or in a

IPR2025-00943
Tesla EX1002 Page 1510

space defined by the set of predicted variables, or in a multi-dimensional space combining system-state/context/action space and said space defined by the set of predicted variables.

9. The automated action-selection system according to any one of claims 2 to 8, wherein the action selector (4) is adapted to select either a random action or that one of the set of candidate actions that the assessment unit (14) evaluates to be expected to yield the greatest improvement in performance of the prediction unit (12).

10. The automated action-selection system according to any one of claims 1 to 9, wherein the prediction unit (12) is adapted to make predictions by applying nearest-neighbours algorithms.

11. An automated prediction-machine-training system comprising an automated action-selection system according to one of claims 1 to 10, wherein the combination of the region identifier (2) and the prediction unit (12) of the automated action-selection system constitute a trainable prediction machine; the automated prediction-machine-training system comprising means for receiving feedback regarding the actual values of the predicted variables (***VAR***) resulting from performance of a given action; and wherein the prediction means (12) is responsive to the feedback data whereby to adapt the prediction function that is applied for candidate actions in the same region of system-state/context/action space as the region containing the combination of the action which produced the feedback data and any system-state/context parameters applicable at the time when said given action was performed.

12. The automated prediction-machine-training system according to claim 11, and comprising:

monitoring means for monitoring the evolution of the respective prediction functions applied by the prediction unit (12) for the different regions of system-state/context/action space; and
an operation-mode setter (16) adapted to change over operation of the prediction-machine-training system (11) from a training mode to a prediction mode when the monitoring means determines that the rate of change of the prediction functions has fallen below a threshold level.

13. An automated action-selection system according to any one of claims 1 to 10, adapted to control the actions performed by a robot or other self-developing device or system, said actions affecting the setup (***S***).

14. An automated action-selection system (21) according to claim 13, and comprising means for supplying the signal indicating the desirability of selecting a given candidate action for performance to a decision-making unit which selects the actions to be performed by said robot or other self-developing device or system.

15. An automated action-selection system (21) according to claim 13 or 14, wherein the action-selection system (21) is separate from, but in communication with the robot or other self-developing device or system.

16. A self developing robot, or other self-developing device or system, comprising an automated action-selection system (21) according to claim 13 or 14.

17. An automated action-selection method adapted to generate signals specifying values for a set of one or more action variables defining an action that can be taken whereby to affect a setup (***S***), the automated action-selection method comprising the steps of:

deceiving signals indicative of the value, at a time *t*, of a set of system-state/context parameters (***SC(t)***) describing the state and/or context of the setup (***S***);
defining a set of regions in a multi-dimensional system-state/context/action space, each dimension of the system-state/context/action space being defined by a respective different parameter or variable of the sets of system-state/context parameters and action variables;
determining a set of candidate actions, each candidate action consisting of a possible set of values for the action variables;
identifying the region in system-state/context/action space containing the combination of a given candidate action with values of any system-state/context parameters at time ***t***;
predicting the value of a set of one or more predicted variables (***VAR***) a predetermined interval after time ***t***, wherein a prediction function applied in the predicting step depends upon the region in system-state/context/action space containing the combination of this given candidate action with any system-state/context parameters at time ***t***;
calculating, for selected candidate actions, a respective indicator of the expected error in the prediction made

IPR2025-00943
Tesla EX1002 Page 1511

in the predicting step for said selected candidate action;
storing indicators of actual prediction errors made in the predicting step for respective candidate actions selected on one or more previous occasions;
evaluating the expected improvement in the performance of the predicting step if a given candidate action is performed, wherein an assessment performed in the evaluating step depends upon the region (***R***) in system-state/context/action space containing the combination of this given candidate action with the values, at time t, of any system-state/context parameters, and the evaluating step is further adapted to evaluate said expected improvement by comparing an indicator of said expected error in prediction with an indicator of the actual prediction error that existed on one or more occasions, previous to time ***t***, when the setup (***S***) had a combination of system-state/context parameters and action variables located in the same region (***R***) of the system-state/context/action space;
generating a signal indicating the desirability of selecting a given candidate action for performance, said signal being dependent on the expected improvement in the performance of the prediction unit (12) evaluated in the evaluating step for said given candidate action;
selecting one of the candidate actions for performance based on said signals indicating the desirability of selecting a given candidate; and
outputting a signal indicating the selected candidate action.

18. The automated action-selection method according to claim 17, and comprising the step of feeding back data indicative of the actual values of the predicted variables (***VAR***) resulting from performance of a given action, and data indicative of the applicable values of any system-state/context parameters at the time said given action was performed, wherein the predicting step is responsive to the feedback data whereby to adapt the prediction function that is applied for candidate actions in the same region of system-state/context/action space as the region containing the combination of the action which produced the feedback data and any system-state/context parameters applicable at the time when said given action was performed.

19. An automated action-selection method according to claim 17 or 18, and comprising the step of controlling the actions performed by a robot or other self-developing device or system, said actions affecting the setup (***S***).

20. An automated action-selection method according to claim 19, and comprising the step of supplying the signal indicating the desirability of selecting a given candidate action for performance to a decision-making unit which selects the actions to be performed by said robot or self-developing device or system.

21. A self-developing robot, or other self-developing device or system, trained by having been caused to participate in the performance of a series of actions, each of the actions in the series having been selected using the automated action-selection system of any one of claims 13 to 15 and/or using the automated action-selection method of claim 19 or 20, wherein the system-state/context parameters (***SC***(***t***)) resulting from performance, by said robot or other self-developing device, of a given action in said series constituted the system-state/context parameters (***SC***(***t***)) used by the automated action-selection system or method in selection of the next action in said series.

22. A prediction machine trained by having been caused to predict a series of sets of predicted variables (VAR) using the automated prediction-machine-training system of claim 11 or 12, data indicative of the actual values of the predicted variables (***VAR***) resulting from performance of a given action, as well as data indicative of the applicable values of any system-state/context parameters at the time said given action was performed, having been fed back to the prediction machine such that an adaptation, responsive to the feedback data, was made of the prediction function applied for candidate actions in the same region of system-state/context/action space as the region containing the combination of the action which produced the feedback data and any system-state/context parameters applicable at the time when said given action was performed.

23. A computer system adapted to predict the value of a set of one or more predicted variables (***VAR***), said computer system having been trained by having been caused to predict a series of sets of predicted variables (VAR) using the automated prediction-machine-training system of claim 11 or 12, data indicative of the actual values of the predicted variables (***VAR***) resulting from performance of a given action, as well as data indicative of the applicable values of any system-state/context parameters at the time said given action was performed, having been fed back to the computer system such that an adaptation, responsive to the feedback data, was made of the prediction function applied for candidate actions in the same region of system-state/context/action space as the region containing the combination of the action which produced the feedback data and any system-state/context parameters applicable at the time when said given action was performed.

IPR2025-00943
Tesla EX1002 Page 1512

24. An expert system adapted to predict the value of a set of one or more predicted variables (***VAR***), said expert system constituting a prediction-machine having been trained by having been caused to predict a series of sets of predicted variables (VAR) using the automated prediction-machine-training system of claim 11 or 12, data indicative of the actual values of the predicted variables (***VAR***) resulting from performance of a given action, as well as data indicative of the applicable values of any system-state/context parameters at the time said given action was performed, having been fed back to the prediction machine such that an adaptation, responsive to the feedback data, was made of the prediction function applied for candidate actions in the same region of system-state/context/action space as the region containing the combination of the action which produced the feedback data and any system-state/context parameters applicable at the time when said given action was performed.

Patentansprüche

1. Automatisiertes Aktionsauswahlsystem (1), welches dazu ausgebildet ist, Signale zu erzeugen, die Werte für einen Satz aus einer oder mehreren Aktionsvariablen spezifizieren, welche eine Aktion definieren, die ausgeführt werden kann, wodurch eine Konfiguration (S) beeinflusst wird,
wobei das automatische Aktionsauswahlsystem (1) aufweist:

eine Eingabeeinrichtung zum Empfangen von Signalen, die zu einem Zeitpunkt t den Wert eines Satzes von Systemzustands-/-kontextparametern (SC(t)) anzeigen, die den Zustand und/oder den Kontext der Konfiguration (S) beschreiben,
einen Bereichsdefinierer (2), welcher dazu ausgebildet ist, einen Satz von Bereichen in einem multidimensionalen Systemzustands-/-kontext-/- aktionsraum zu definieren, wobei jede Dimension des Systemzustands-/-kontext-/-aktionsraums definiert wird durch einen jeweiligen unterschiedlichen Parameter oder eine Variable aus den Sätzen von Systemzustands-/- kontextparametern und Aktionsvariablen,
eine Einrichtung zum Bestimmen eines Satz von Kandidatenaktionen, wobei jede Kandidatenaktion aus einem möglichen Satz von Werten für die Aktionsvariablen besteht,
einen Bereichsidentifizierer (2) zum Identifizieren des Bereichs im Systemzustands-/-kontext-/-aktionsraum, welcher die Kombination aus einer gegebenen Kandidatenaktion mit Werten eines jeglichen Systemzustands-/-kontextparameters zu einem Zeitpunkt t aufweist,
eine Vorhersageeinheit (12), welche dazu ausgebildet ist, den Wert eines Satzes einer oder mehrerer vorhergesagter Variablen (VAR) eine vorbestimmte Zeitspanne nach der Zeit t vorherzusagen, wobei eine Vorhersagefunktion, die durch die Vorhersageeinheit (12) angewandt wird, abhängt von dem Bereich im Systemzustands-/-köntext-/-aktionsraum, welcher die Kombination dieser gegebenen Kandidatenaktion mit jeglichen Systemzustands-/-kontextparametern zum Zeitpunkt t aufweist,
einer Berechnereinrichtung, welche dazu ausgebildet ist, für ausgewählte Kandidatenaktionen einen jeweiligen Indikator für den erwarteten Fehler in der Vorhersage zu berechnen, die von der Vorhersageeinheit (12) für die ausgewählte Kandidatenaktion gemacht wird,
eine Speichereinrichtung zum Speichern von Indikatoren tatsächlicher Vorhersagefehler, die von der Vorhersageeinheit (12) für jeweilige Kandidatenaktionen gemacht werden, die zu einer oder mehreren vorangehenden Gelegenheiten ausgewählt wurden,
eine Prüfeinrichtung (14), welche dazu ausgebildet ist, die erwartete Verbesserung in der Performance der Vorhersageeinheit (12) zu bewerten, falls eine gegebene Kandidatenaktion ausgeführt wird, wobei eine Prüfung, die von der Prüfeinrichtung (14) ausgeführt wird, von dem Bereich (R) im Systemzustands-/-kontext-/-aktionsraum abhängt, welcher die Kombination dieser gegebenen Kandidatenaktion mit den Werten eines jeglichen Systemzustands-/-kontextparameters zum Zeitpunkt t enthält, und wobei die Prüfeinrichtung des Weiteren dazu ausgebildet ist, die erwartete Verbesserung zu bewerten durch Vergleichen eines Indikators des erwarteten Fehlers in der Vorhersage mit einem Indikator des tatsächlichen Vorhersagefehlers, der zu einer oder mehreren Gelegenheiten vor dem Zeitpunkt t existierte, wenn die Konfiguration (S) eine Kombination von Systemzustands-/- kontextparametern und Aktionsvariablen aufwies, die im selben Bereich (R) des Systemzustands-/-kontext-/-aktionsraums lokalisiert war,
eine Einrichtung zum Erzeugen eines Signals, welches die Erwünschtheit des Auswählens einer gegebenen Kandidatenaktion für die Performance anzeigt, wobei das Signal abhängig ist von der erwarteten Verbesserung in der Performance der Vorhersageeinheit (12), welche durch die Prüfeinheit (14) für die gegebene Kandidatenaktion bewertet wurde, und
eine Aktionsauswahleinrichtung (4), welche dazu ausgebildet ist, eine der Kandidatenaktionen für die Performance auf der Grundlage der Signale auszuwählen, welche die Erwünschtheit des Auswählens einer gegebenen Kandidatenaktion anzeigen, und welche des Weiteren dazu ausgebildet ist, ein Signal, welches die ausgewählte

IPR2025-00943
Tesla EX1002 Page 1513

Kandidatenaktion anzeigt, auszugeben.

2. Automatisiertes Aktionsauswahlsystem nach Anspruch 1,
welches des Weiteren eine Aktionsauswahleinrichtung (4) aufweist, die dazu ausgebildet ist, eine Aktion für die Performance zum Beeinflussen der Konfiguration (S) auszuwählen, wobei die Aktionsauswahleinrichtung (4) dazu ausgebildet ist, eine Wahrscheinlichkeit p, wobei $0 < p < 1$ gilt, für das Auswählen derjenigen aus dem Satz von Kandidatenaktionen hat, welche die Prüfeinheit (14) auswertet, von der erwartet wird, dass sie die größte Verbesserung in der Performance der Vorhersageeinheit (12) erzielt, wobei das Aktionsauswahlsystem (1) dazu ausgebildet ist, diejenigen Daten auszugeben, die die von der Aktionsauswahleinrichtung (4) ausgewählte Aktion definieren.

3. Automatisiertes Aktionsauswahlsystem nach Anspruch 1 oder 2,
wobei der Bereichsdefinierer (2) dazu ausgebildet ist, einen existierenden Bereich (R) in zwei oder mehr neue Bereiche zu unterteilen, wenn ein erstes Kriterium (C1) erfüllt ist.

4. Automatisiertes Aktionsauswahlsystem nach Anspruch 3,
welche des Weiteren eine Zähleinrichtung aufweist zum Zählen der Anzahl von Gelegenheiten, bei welchen eine erste und eine zweite Bedingung beide erfüllt sind, wobei die erste Bedingung darin besteht, dass eine Aktion, die ausgeführt wird, die Konfiguration (S) beeinflusst und wobei die zweite Bedingung darin besteht, dass die Kombination der Aktionsvariablenwerte, welche die Aktion definieren, und der Systemzustands-/-kontextparameterwerte zu dem Zeitpunkt, bei welchem die Aktion ausgeführt wird, in den Bereich (R) fällt, und
wobei das erste Kriterium (C1) erfüllt ist, wenn die Zähleinrichtung bis zu einer vorbestimmten Zahl (NS) hochgezählt hat.

5. Automatisiertes Aktionsauswahlsystem nach einem der Ansprüche 1 bis 4,
wobei des Weiteren eine Metavorhersageeinrichtung (13) vorgesehen ist, welche dazu ausgebildet ist, den erwarteten Fehler in den Werten des Satzes von Variablen (VAR) zu bewerten, die durch die Vorhersageeinrichtung (12) vorhergesagt wurden, wobei eine Vorhersagefunktion, die von der Metavorhersageeinrichtung (12) angewandt wird, abhängig ist vom Bereich im Systemzustands-/-kontext-/-aktionsraum, welcher die Kombination der gegebenen Kandidatenaktion mit den Systemzustands-/-kontextparametern zum Zeitpunkt t enthält, wobei die Prüfeinheit (14) dazu ausgebildet ist, die erwartete Verbesserung in der Performance der Vorhersageeinrichtung (12) zu bewerten, wenn eine gegebene Kandidatenaktion ausgeführt wird, durch Auswerten des Abfalls im Vorhersagefehler, von dem erwartet wird, dass er aus der Performance der gegebenen Kandidatenaktion resultiert.

6. Automatisiertes Aktionsauswahlsystem nach einem der Ansprüche 1 bis 5,
wobei eine Einrichtung vorgesehen ist zum Empfangen einer Rückmeldung in Bezug auf die tatsächlichen Werte der vorhergesagter Variablen (VAR), die aus der Performance einer gegebenen Aktion resultieren,
wobei die Vorhersageeinrichtung (12) auf die Rückmeldedaten reagierend ist, um die Vorhersagefunktion anzupassen, die für Kandidatenaktionen im selben Bereich des Systemzustands-/-kontext-/-aktionsraums angewandt wird wie im Bereich, welcher die Kombination enthält aus der Aktion, welche die Rückmeldungsdaten erzeugte, und den Werten eines jeglichen Systemzustands-/-kontextparameters zu dem Zeitpunkt, bei dem die gegebene Aktion ausgeführt wurde.

7. Automatisiertes Aktionsauswahlsystem nach Anspruch 6,
welches eine Speichereinrichtung aufweist zum Speichern von Trainingsbeispieldaten,
wobei für eine gegebene ausgeführte Aktion die Trainingsbeispieldaten Aktionsvariablenwerte, jegliche Systemzustands-/-kontextparameterwerte, die zu dem Zeitpunkt anwendbar waren, zu dem die gegebene Aktion ausgeführt wurde, und Rückmeldungsdaten, welche die tatsächlichen Werte des Satzes vorhergesagter Variablen (VAR) definieren, aufweisen.

8. Automatisiertes Aktionsauswahlsystem nach Anspruch 7,
wobei der Bereichsdefinierer (2) dazu ausgebildet ist, einen Bereich (R) des Systemzustands-/-kontext-/-aktionsraums in zwei oder mehr neue Bereiche zu unterteilen, wenn ein erstes Kriterium (C1) erfüllt ist, durch Definieren von Grenzen der neuen Bereiche derart, dass ein Gleichgewicht vorliegt zwischen den Anzahlen der Trainingsbeispiele in jedem neuen Bereich und dass die Varianz der Trainingsbeispiele im Systemzustands-/-kontext-/- aktionsraum oder in einem Raum, der definiert wird durch den Satz vorhergesagter Variablen oder in einem multidimensionalen Raum, der den Systemzustands-/-kontext-/-aktionsraum und den durch den Satz vorhergesagter Variablen definierten Raum miteinander kombiniert, minimiert wird.

IPR2025-00943
Tesla EX1002 Page 1514

9. Automatisiertes Aktionsauswahlsystem nach einem der Ansprüche 2 bis 8,
wobei die Aktionsauswahleinrichtung (4) dazu ausgebildet ist, entweder eine zufällige Aktion oder diejenige aus dem Satz von Kandidatenaktionen auszuwählen, welche die Prüfeinheit (14) auswertet, von der erwartet wird, dass sie die größte Verbesserung in der Performance der Vorhersageeinheit (12) liefert.

10. Automatisiertes Aktionsauswahlsystem nach einem der Ansprüche 1 bis 9,
wobei die Vorhersageeinheit (12) dazu ausgebildet ist, Vorhersagen durch Anwenden von Nächstnachbaralgorithmen zu machen.

11. Automatisiertes Vorhersagemaschinentrainingssystem mit einem automatisierten Aktionsauswahlsystem nach einem der Ansprüche 1 bis 10,
wobei die Kombination des Bereichsidentifizierers (2) und der Vorhersageeinheit (12) des automatisierten Aktionsauswahlsystems eine trainierbare Vorhersagemaschine bildet,
wobei das automatisierte Vorhersagemaschinentrainingssystem eine Einrichtung aufweist zum Empfangen einer Rückmeldung in Bezug auf die tatsächlichen Werte der vorhergesagten Variablen (VAR), welche aus der Performance einer gegebenen Aktion resultieren, und
wobei die Vorhersageeinrichtung (12) auf die Rückmeldedaten reagierend ist, um die Vorhersagefunktion anzupassen, die für Kandidatenaktionen im selben Bereich des Systemzustands-/-kontext-/-aktionsraums angewandt wird wie im Bereich, welcher die Kombination enthält aus der Aktion, welche die Rückmeldungsdaten erzeugte, und jeglichen Systemzustands-/- kontextparametern, die anwendbar waren zum Zeitpunkt, bei welchem die gegebene Aktion ausgeführt wurde.

12. Automatisiertes Vorhersagemaschinentrainingssystem nach Anspruch 11,
welches eine Überwachungseinrichtung aufweist zum Überwachen der Entwicklung der jeweiligen Vorhersagefunktionen, die von der Vorhersageeinheit (12) für die verschiedenen Bereichen des Systemzustands-/-kontext/-aktionsraums angewandt werden, und
mit einer Betriebsmoduseinstelleinrichtung (16), welche dazu ausgebildet ist, den Betrieb des Vorhersagemaschinentrainingssystems (11) zu wechseln von einem Trainingsmodus zu einem Vorhersagemodus, wenn die Überwachungseinrichtung ermittelt, dass die Änderungsrate der Vorhersagefunktionen unterhalb eines Schwellwerteniveaus gefallen ist.

13. Automatisiertes Aktionsauswahlsystem nach einem der Ansprüche 1 bis 10,
welches dazu ausgebildet ist, die Aktionen zu steuern, die von einem Roboter oder von einer anderen selbstentwickelnden Einrichtung oder System ausgeführt werden und welche die Konfiguration (S) beeinflussen.

14. Automatisiertes Aktionsauswahlsystem (21) nach Anspruch 13,
mit einer Einrichtung zum Zuführen des Signals, welches die Erwünschtheit des Auswählens einer gegebenen Kandidatenaktion für die Performance einer Entscheidungseinheit anzeigt, welche die von dem Roboter oder der anderen selbstentwickelnden Einrichtung oder System auszuführenden Aktionen auswählt.

15. Automatisiertes Aktionsauswahlsystem (21) nach Anspruch 13 oder 14,
wobei das Aktionsauswahlsystem (21) getrennt ist von dem aber in Kommunikation steht mit dem Roboter oder der anderen selbstentwickelnden Einrichtung oder System.

16. Selbstentwickelnder Roboter oder andere selbstentwickelnde Einrichtung oder System,
mit einem automatisierten Aktionsauswahlsystem (21) nach Anspruch 13 oder 14.

17. Automatisiertes Aktionsauswahlverfahren, welches dazu ausgebildet ist, Signale zu erzeugen, welche Werte spezifizieren für einen Satz einer oder mehrerer Aktionsvariablen, welche eine Aktion definieren, die ausgeführt werden kann, wodurch eine Konfiguration (S) beeinflusst wird,
wobei das automatisierte Aktionsauswahlverfahren die Schritte aufweist:

Empfangen von Signalen, die zu einem Zeitpunkt t anzeigend sind für den Wert eines Satzes von Systemzustands-/-kontextparametern (SC(t)), welche den Zustand und/oder den Kontext der Konfiguration (S) beschreiben,
Definieren eines Satzes von Bereichen in einem multidimensionalen Systemzustands-/-kontext-/-aktionsraum, wobei jede Dimension des Systemzustands-/-kontext-/-aktionsraums definiert wird durch einen jeweils unterschiedlichen Parameter oder eine Variable der Sätze von Systemzustands-/-kontextparametern und Aktions-

IPR2025-00943
Tesla EX1002 Page 1515

variablen,
Bestimmen eines Satzes von Kandidatenaktionen, wobei jede Kandidatenaktion aus einem möglichen Satz von Werten für die Aktionsvariablen besteht,
Identifizieren des Bereichs im Sysiemzustands-/-kontext-/- aktionsraum, welcher die Kombination einer gegebenen Kandidatenaktion mit Werten eines jeglichen Systemzustands-/-kontextparameters zu einem Zeitpunkt t enthält,
Vorhersagen des Werts eines Satzes einer oder mehrerer vorhergesagter Variablen (VAR) eine vorbestimmte Zeitspanne nach der Zeit t, wobei eine Vorhersagefunktion, die im Vorhersageschritt angewandt wird, abhängt von dem Bereich im Systemzustands-/-kontext-/-aktionsraum, welcher die Kombination dieser gegebenen Kandidatenaktion mit einem jeglichen Systemzustands-/-kontextparameter zum Zeitpunkt t enthält,
für ausgewählte Kandidatenaktionen Berechnen eines jeweiligen Indikators des erwarteten Fehlers in der Vorhersage, der im Vorhersageschritt für die ausgewählte Kandidatenaktion gemacht wurde,
Speichern von Indikatoren tatsächlicher Vorhersagefehler, die im Vorhersageschritt für jeweilige Kandidatenaktionen, die bei einer oder mehreren vorangehenden Gelegenheiten ausgewählt wurden,
Bewerten der erwarteten Verbesserung in der Performance des Vorhersageschritts, falls eine gegebene Kandidatenaktion ausgeführt wird, wobei eine Überprüfung, die im Auswerteschritt ausgeführt wurde, abhängig ist von dem Bereich (R) im Systemzustands-/-kontext-/-aktionsraum, welcher die Kombination dieser gegebenen Kandidatenaktion mit den Werten eines jeglichen Systemzustands-/-kontextparameters zum Zeitpunkt t enthält, und wobei der Bewertungsschritt des Weiteren ausgebildet ist, um die erwartete Verbesserung zu bewerten durch Vergleichen eines Indikators des erwarteten Fehlers in der Vorhersage mit einem Indikator des tatsächlichen Vorhersagefehlers, der bei einer oder mehreren Gelegenheiten vor dem Zeitpunkt t existierte, wenn die Konfiguration (S) eine Kombination von Systemzustands-/-kontextparametern und Aktionsvariablen hatte, die im selben Bereich (R) des Systemzustands-/-kontext-/-aktionsraums lokalisiert war,
Erzeugen eines Signals, welches die Erwünschtheit des Auswählens einer gegebenen Kandidatenaktion für die Performance anzeigt, wobei das Signal abhängig ist von der erwarteten Verbesserung in der Performance der Vorhersageeinheit (12), die im Bewertungsschritt für die gegebene Kandidatenaktion bewertet wurde,
Auswählen einer der Kandidatenaktionen für die Performance auf der Grundlage der Signale, die die Erwünschtheit des Auswählens eines gegebenen Kandidaten anzeigen, und
Ausgeben eines Signals, welches die ausgewählte Kandidatenaktion anzeigt.

18. Automatisiertes Aktionsauswahlverfahren nach Anspruch 17,
mit einem Schritt des Rückmeldens von Daten, die anzeigend sind für die tatsächlichen Werte der vorhergesagten Variablen (VAR), welche aus der Performance einer gegebenen Aktion resultieren, und von Daten, die anzeigend sind für die anwendbaren Werte eines jeglichen Systemzustands-/- kontextparameters zum Zeitpunkt, bei welchem die gegebene Aktion ausgeführt wurde, wobei der Vorhersageschritt reagierend ist auf die Rückmeldedaten, wodurch die Vorhersagefunktion angepasst wird, welche angewandt wird für die Kandidatenaktion im selben Bereich des Systemzustands-/- kontext-/-aktionsraums wie dem Bereich, welcher die Kombination der Aktion, welche die Rückmeldungsdaten erzeugte, und jegliche Systemzustands/-kontextparamtern enthält, welche zu dem Zeitpunkt anwendbar waren, als die gegebene Aktion ausgeführt wurde.

19. Automatisiertes Aktionsauswahlverfahren nach Anspruch 17 oder 18,
mit dem Schritt des Steuerns der Aktionen, die ausgeführt werden durch einen Roboter oder eine andere selbstentwickelnde Einrichtung oder ein System, wobei die Aktionen die Konfiguration (S) beeinflussen.

20. Automatisiertes Aktionsauswahlverfahren nach Anspruch 19,
mit dem Schritt des Zuführens des Signals, welches die Erwünschtheit des Auswählens einer gegebenen Kandidatenaktion für die Performance an eine Entscheidungseinheit, welche die Aktionen auswählt, die durch den Roboter oder die selbstentwickelnde Einrichtung oder das System auszuführen sind.

21. Selbstentwickelnder Roboter oder eine andere selbstentwickelnde Einrichtung oder System,
welche trainiert wurden, indem bewirkt wurde, dass sie in der Performance einer Serie von Aktionen teilnehmen, wobei jede der Aktionen in der Serie ausgewählt wurde durch Verwenden des automatisierten Aktionsauswahlsystems nach einem der Ansprüche 13 bis 15 und/oder durch Verwenden des automatisierten Aktionsauswahlverfahrens nach Anspruch 19 oder 20,
wobei die Systemzustands-/-kontextparameter (SC(t)), die aus der Performance in Bezug auf eine gegebene Aktion aus der Serie durch den Roboter oder die andere selbstentwickelnde Einrichtung resultieren, die Systemzustands-/-kontextparameter (SC(t)) bilden, die vom automatisierten Aktionsauswahlsystem oder dem Verfahren bei der Auswahl der nächsten Aktion in der Serie verwendet werden.

IPR2025-00943
Tesla EX1002 Page 1516

22. Vorhersagemaschine, die trainiert wurde, indem bewirkt wurde, dass sie eine Serie von Sätzen von vorhergesagten Variablen (VAR) unter Verwendung des automatisierten Vorhersagemaschinentrainingssystems nach Anspruch 11 oder 12 vorhersagt,
wobei Daten, die anzeigend sind für die tatsächlichen Werte der vorhergesagten Variablen (VAR), die aus der Performance einer gegebenen Aktion resultieren, und Daten, die anzeigend sind für die anwendbaren Werte eines jeweiligen Systemzustands-/-kontextparameters zu dem Zeitpunkt, bei welchem die gegebene Aktion ausgeführt wurde, an die Vorhersagemaschine derart zurück gemeldet wurden, dass eine Adaption in Antwort auf die Rückmeldedaten für die Vorhersagefunktion, welche für Kandidatenaktionen angewandt wurde im selben Bereich des Systemzustands-/-kontext-/- aktionsraums wie dem Bereich, der die Kombination der Aktion, welche die Rückmeldungsdaten erzeugte, und jegliche Systemzustands-/- kontextparameter aufweist, die zu dem Zeitpunkt anwendbar sind, bei dem die gegebene Aktion ausgeführt wurde, durchgeführt wurde.

23. Computersystem, welches dazu ausgebildet ist, den Wert eines Satzes aus einer oder mehreren vorhergesagten Variablen (VAR) vorherzusagen, wobei das Computersystem trainiert wurde, indem bewirkt wurde, dass sie eine Serie von Sätzen von vorhergesagten Variablen (VAR) unter Verwendung des automatisierten Vorhersagemaschinentrainingssystems nach Anspruch 11 oder 12 vorhersagt,
wobei Daten, die anzeigend sind für die tatsächlichen Werte der vorhergesagten Variablen (VAR), die aus der Performance einer gegebenen Aktion resultieren, und Daten, die anzeigend sind für die anwendbaren Werte eines jeglichen Systemzustands-/-kontextparameters zu dem Zeitpunkt, bei welchem die gegebene Aktion ausgeführt wurde, an das Computersystem derart zurück gemeldet wurden, dass eine Adaption in Antwort auf die Rückmeldedaten in Bezug auf die Vorhersagefunktion, welche für Kandidatenaktionen angewandt wurde im selben Bereich des Systemzustands-/- kontext-/-aktionsraunis wie dem Bereich, der die Kombination der Aktion, welche die Rückmeldungsdaten erzeugte, und jegliche Systemzustands-/- kontextparameter aufweist, die zu dem Zeitpunkt anwendbar sind, bei dem die gegebene Aktion ausgeführt wurde, durchgeführt wurde.

24. Expertensystem, die trainiert wurde, indem bewirkt wurde, dass sie eine Serie von Sätzen von vorhergesagten Variablen (VAR) unter Verwendung des automatisierten Vorhersagemaschinentrainingssystems nach Anspruch 11 oder 12 vorhersagt,
wobei Daten, die anzeigend sind für die tatsächlichen Werte der vorhergesagten Variablen (VAR), die aus der Performance einer gegebenen Aktion resultieren, und Daten, die anzeigend sind für die anwendbaren Werte eines jeweiligen Systemzustands-/-kontextparameters zum Zeitpunkt, bei welchem die gegebene Aktion ausgeführt wurde, an die Vorhersagemaschine derart zurück gemeldet wurden, dass eine Adaption in Antwort auf die Rückmeldedaten in Bezug auf die Vorhersagefunktion, welches für Kandidatenaktionen angewandt wurde im selben Bereich des Systemzustands-/- kontext-/-aktionsraums als den Bereich, der die Kombination der Aktion, welche die Rückmeldungsdaten erzeugt, und jeglichen Systemzustands-/- kontextparameters aufweist, der zu dem Zeitpunkt anwendbar ist, bei dem die gegebene Aktion ausgeführt wurde, durchgeführt wurde.

Revendications

1. Système de sélection d'action automatisée (1) conçu pour générer des signaux spécifiant des valeurs pour un ensemble d'une ou de plusieurs variables d'action définissant une action qui peut être prise pour affecter une configuration (S), le système de sélection d'action automatisée (1) comprenant :

un moyen d'entrée destiné à recevoir des signaux indicatifs de la valeur, à un temps t, d'un ensemble de paramètres d'état-de-système/contexte (SC(t)) décrivant l'état et/ou le contexte de la configuration (S) ;
un définisseur de région (2) conçu pour définir un ensemble de régions dans un espace d'état-de-système/contexte/action multidimensionnel, chaque dimension de l'espace d'état-de-système/contexte/action étant défini par une variable ou un paramètre différent respectif des ensembles des variables d'action et paramètres d'état-de-système/contexte;
un moyen destiné à déterminer un ensemble d'actions candidates, chaque action candidate se composant d'un ensemble possible de valeurs pour les variables d'action ;
un identificateur de région (2) destiné à identifier la région dans un espace d'état-de-système/contexte/action contenant la combinaison d'une action candidate donnée avec des valeurs de tout paramètre d'état-de-système/contexte au temps t ;
une unité de prévision (12) conçue pour prévoir la valeur d'un ensemble d'une ou de plusieurs variables prévues (VAR) à un intervalle prédéterminé après le temps t, dans lequel une fonction de prévision appliquée par l'unité de prévision (12) dépend de la région de l'espace d'état-de-système/contexte/action contenant la combinaison

IPR2025-00943
Tesla EX1002 Page 1517

de cette action candidate donnée avec tout paramètre d'état-de-système/contexte au temps t ;
un moyen de calcul conçu pour calculer, pour les actions candidates sélectionnées, un indicateur respectif de l'erreur attendue lors de la prévision effectuée par l'unité de prévision (12) pour ladite action candidate sélectionnée ;
un moyen de mémoire destiné à stocker des indicateurs d'erreurs de prévision réelles réalisées par l'unité de prévision (12) pour les actions candidates respectives sélectionnées à une ou plusieurs occasions précédentes ;
un moyen d'évaluation (14) conçu pour évaluer l'amélioration attendue sur les performances de l'unité de prévision (12) si une action candidate donnée est réalisée, dans lequel une évaluation effectuée par le moyen d'évaluation (14) dépend de la région (R) dans l'espace d'état-de-système/contexte/action contenant la combinaison de cette action candidate donnée avec les valeurs, au temps t, de toute paramètre d'état-de-système/contexte, et le moyen d'évaluation est en outre conçu pour évaluer ladite amélioration attendue en comparant un indicateur de ladite erreur de prévision attendue avec un indicateur de l'erreur de prévision réelle qui existait à une ou plusieurs occasions, avant le temps t, lorsque la configuration (S) avait une combinaison de paramètres d'état-de-système/contexte et de variables d'action se trouvant dans la même région (R) de l'espace d'état-de-système/contexte/action;
un moyen destiné à générer un signal indiquant les avantages de sélectionner une action candidate donnée en vue de la réalisation de celle-ci, ledit signal dépendant de l'amélioration attendue sur les performances de l'unité de prévision (12) évaluée par l'unité d'évaluation (14) pour ladite action candidate donnée ; et
un sélecteur d'action (4) conçu pour sélectionner l'une des actions candidates, en vue de la réalisation de celle-ci, sur la base desdits signaux indiquant les avantages de sélectionner une action candidate donnée et conçu en outre pour émettre en sortie un signal indiquant l'action candidate sélectionnée.

2. Système de sélection d'action automatisée selon la revendication 1, et comprenant un sélecteur d'action (4) conçu pour sélectionner une action en vue de la réalisation de celle-ci afin d'affecter la configuration (S), le sélecteur d'action (4) est conçu pour avoir une probabilité de p, où $0<p<1$, de sélectionner l'action de l'ensemble d'actions candidates que l'unité d'évaluation (14) évalue comme étant attendue afin de produire la meilleure amélioration en matière de performances de l'unité de prévision (12) ; dans lequel le système de sélection d'action (1) est conçu pour émettre en sortie des données définissant l'action sélectionnée par le sélecteur d'action (4).

3. Système de sélection d'action automatisée selon la revendication 1 ou 2, dans lequel le définisseur de région (2) est conçu pour diviser une région existante (R) en deux, ou plus, nouvelles régions lorsqu'un premier critère 9C1) est rencontré.

4. Système de sélection d'action automatisée selon la revendication 3, et comprenant un moyen de comptage destiné à compter le nombre d'occasions lors desquelles une première et une seconde condition sont toutes les deux rencontrées, la première condition étant qu'une action est prise, affectant la configuration (S), et la seconde condition étant que la combinaison des valeurs de variable d'action définissant ladite action et des valeurs de paramètre d'état-de-système/contexte au moment où ladite action est prise tombe dans la région (R) ;
dans lequel le premier critère (C1) est rencontré lorsque le moyen de comptage a compté jusqu'à un nombre prédéterminé (NS).

5. Système de sélection d'action automatisée selon l'une quelconque des revendications 1 à 4 et comprenant un moyen de méta-prévision (13) conçu pour évaluer l'erreur attendue sur les valeurs dudit ensemble de variables (VAR) prévues par le moyen de prévision (12), une fonction de prévision appliquée par le moyen de méta-prévision (13) dépendant de la région de l'espace d'état-de-système/contexte/action contenant la combinaison de ladite action candidate donnée avec les paramètres d'état-de-système/contexte au temps t ; dans lequel le moyen d'évaluation (14) est conçu pour évaluer l'amélioration attendue en matière de performances du moyen de prévision (12) si une action candidate donnée est réalisée en évaluant la diminution d'erreur de prévision qui est attendue en relation avec la réalisation de ladite action candidate donnée.

6. Système de sélection d'action automatisée selon l'une quelconque des revendications 1 à 5, et comprenant un moyen destiné à recevoir un retour concernant les valeurs réelles des variables prévues (VAR) résultantes de la réalisation d'une action donnée ;
le moyen de prévision (12) répond aux données de retour pour adapter la fonction de prévision qui est appliquée pour les actions candidates de la même région de l'espace d'état-de-système/contexte/action que la région contenant la combinaison de l'action qui a produit les données de retour et les valeurs de tout paramètre d'état-de-système/contexte au moment où ladite action donnée a été réalisée.

IPR2025-00943
Tesla EX1002 Page 1518

7. Système de sélection d'action automatisée selon la revendication 6, et comprenant un moyen de mémoire destiné à stocker des données d'exemples de formation ; pour une action réalisée donnée, lesdits données d'exemples de formation comprenant des valeurs de variable d'action, toute valeur de paramètre d'état-de-système/contexte applicable au moment où ladite action donnée a été réalisée, et les données de retour définissant les valeurs réelles de l'ensemble de variables prévues (VAR).

8. Système de sélection d'action automatisée selon la revendication 7, dans lequel le définisseur de région (2) est conçu pour diviser une région (R) de l'espace d'état-de-système/contexte/action en deux, ou plus, nouvelles régions, lorsqu'un premier critère (C1) est rencontré, en définissant les limites des nouvelles régions de sorte qu'il y ait un équilibre entre les nombres d'exemples de formation dans chaque nouvelle région, et que la variance des exemples de formation soit minimisée dans l'espace d'état-de-système/contexte/action, ou dans un espace défini par l'ensemble de variables prévues, ou dans un espace multidimensionnel combinant l'espace d'état-de-système/contexte/action et ledit espace défini par l'ensemble de variables prévues.

9. Système de sélection d'action automatisée selon l'une quelconque des revendications 2 à 8, dans lequel le sélecteur d'action (4) est conçu pour sélectionner une action aléatoire ou l'action de l'ensemble d'actions candidates que l'unité d'évaluation (14) évalue comme étant attendue pour produire la meilleure amélioration en matières de performances de l'unité de prévision (12).

10. Système de sélection d'action automatisée selon l'une quelconque des revendications 1 à 9, dans lequel l'unité de prévision (12) est conçue pour effectuer des prévisions en appliquant des algorithmes du plus proche voisin.

11. Système de formation de machine de prévision automatisée comprenant un système de sélection d'action automatisée selon l'une des revendications 1 à 10, dans lequel la combinaison de l'identificateur de région (2) et de l'unité de prévision (12) du système de sélection d'action automatisée constitue une machine de prévision pouvant être formée ; le système de formation de machine de prévision automatisée comprenant un moyen destiné à recevoir un retour concernant les valeurs réelles des variables prévues (VAR) résultant de la réalisation d'une action donnée ; et dans lequel le moyen de prévision (12) répond aux données de retour pour adapter la fonction de prévision qui est appliquée pour les actions candidates de la même région de l'espace d'état-de-système/contexte/action que la région contenant la combinaison de l'action qui a produit les données de retour et tout paramètre d'état-de-système/contexte applicable au moment où ladite action donnée a été réalisée.

12. Système de formation de machine de prévision automatisée selon la revendication 11, et comprenant :

    un moyen de contrôle destiné à contrôler l'évolution des fonctions de prévision respectives appliquées par l'unité de prévision (12) pour les différentes régions de l'espace d'état-de-système/contexte/action; et
    un configurateur de mode de fonctionnement (16) conçu pour faire changer le fonctionnement du système de formation de machine de prévision (11) en le faisant passer d'un mode formation à un mode prévision lorsque le moyen de contrôle détermine que le taux de changement des fonctions de prévision est tombé sous un niveau seuil.

13. Système de sélection d'action automatisée selon l'une quelconque des revendications 1 à 10, conçu pour commander les actions réalisées par un robot ou autre dispositif ou système d'auto-développement, lesdites actions affectant la configuration (S).

14. Système de sélection d'action automatisée (21) selon la revendication 13, et comprenant un moyen destiné à envoyer le signal indiquant les avantages de sélectionner une action candidate donnée en vue de la réalisation de celle-ci vers une unité de prise d'action qui sélectionne les actions à réaliser par ledit robot ou autre dispositif ou système d'auto-développement.

15. Système de sélection d'action automatisée (21) selon la revendication 13 ou 14, dans lequel le système de sélection d'action (21) est séparé du robot ou autre dispositif ou système d'auto-développement, mais en communication avec celui-ci.

16. Robot d'auto-développement, ou autre dispositif ou système d'auto-développement, comprenant un système de sélection d'action automatisée (21) selon la revendication 13 ou 14.

17. Procédé de sélection d'action automatisée conçu pour générer des signaux spécifiant les valeurs d'un ensemble

IPR2025-00943
Tesla EX1002 Page 1519

d'une ou de plusieurs variables d'action définissant une action qui peut être prise pour affecter une configuration (S), le procédé de sélection d'action automatisée comprenant les étapes consistant à :

recevoir des signaux indicatifs de la valeur, à un temps t, d'un ensemble de paramètres d'état-de-système/contexte (SC(t)) décrivant l'état et/ou le contexte de la configuration (S) ;
définir un ensemble de régions dans un espace d'état-de-système/contexte/action multidimensionnel, chaque dimension de l'espace d'état-de-système/contexte/action étant définie par une variable ou paramètre différent respectif des ensembles de variables d'action et paramètres d'état-de-système/contexte;
déterminer un ensemble d'actions candidates, chaque action candidate se composant d'un ensemble possible de valeurs pour les variables d'action ;
identifier la région dans l'espace d'état-de-système/contexte/action contenant la combinaison d'une action candidate donnée avec des valeurs de tout paramètre d'état-de-système/contexte au temps t ;
prévoir la valeur d'un ensemble d'une ou de plusieurs variables prévues (VAR) à un intervalle prédéterminé après le temps t, dans lequel une fonction de prévision appliquée à l'étape de prévision dépend de la région de l'espace d'état-de-système/contexte/action contenant la combinaison de cette action candidate donnée avec tout paramètre d'état-de-système/contexte au temps t ;
calculer, pour les actions candidates sélectionnées, un indicateur respectif de l'erreur attendue lors de la prévision effectuée à l'étape de prévision pour ladite action candidate sélectionnée ;
stocker des indicateurs d'erreurs de prévision réelles réalisées à l'étape de prévision pour les actions candidates respectives sélectionnées à une ou plusieurs occasions précédentes ;
évaluer l'amélioration attendue sur les performances de l'étape de prévision si une action candidate donnée est réalisée, dans lequel une évaluation effectuée à l'étape d'évaluation dépend de la région (R) dans l'espace d'état-de-système/contexte/action contenant la combinaison de cette action candidate donnée avec les valeurs, au temps t, de toute paramètre d'état-de-système/contexte, et l'étape d'évaluation est en outre conçue pour évaluer ladite amélioration attendue en comparant un indicateur de ladite erreur de prévision attendue avec un indicateur de l'erreur de prévision réelle qui existait à une ou plusieurs occasions, avant le temps t, lorsque la configuration (S) avait une combinaison de paramètres d'état-de-système/contexte et de variables d'action se trouvant dans la même région (R) de l'espace d'état-de-système/contexte/action;
générer un signal indiquant les avantages de sélectionner une action candidate donnée en vue de la réalisation de celle-ci, ledit signal dépendant de l'amélioration attendue sur les performances de l'unité de prévision (12) évaluée par l'étape d'évaluation pour ladite action candidate donnée ;
sélectionner l'une des actions candidates, en vue de la réalisation de celle-ci, sur la base desdits signaux indiquant les avantages de sélectionner une action candidate donnée ;
émettre en sortie un signal indiquant l'action candidate sélectionnée.

18. Procédé de sélection d'action automatisée selon la revendication 17, et comprenant l'étape consistant à retourner des données indicatives des valeurs réelles des variables prévues (VAR) résultant de la réalisation d'une action donnée, et des données indicatives des valeurs applicables de tout paramètre d'état-de-système/contexte au moment où ladite action donnée a été réalisée, dans lequel l'étape de prévision répond aux données de retour pour adapter la fonction de prévision qui est appliquée pour les actions candidates de la même région de l'espace d'état-de-système/contexte/action que la région contenant la combinaison de l'action qui a produit les données de retour et tout paramètre d'état-de-système/contexte applicable au moment où ladite action donnée a été réalisée.

19. Procédé de sélection d'action automatisée selon la revendication 17 ou 18, et comprenant l'étape consistant à commander les actions réalisées par un robot ou autre dispositif ou système d'auto-développement, lesdites actions affectant la configuration (S).

20. Procédé de sélection d'action automatisée selon la revendication 19, et comprenant l'étape consistant à envoyer le signal indiquant les avantages de sélectionner une action candidate donnée en vue de la réalisation de celle-ci vers une unité de prise de décision qui sélectionne les actions à réaliser par ledit robot ou dispositif ou système d'auto-développement.

21. Robot d'auto-développement, ou autre dispositif ou système d'auto-développement, formé en ayant été amené à contribuer à la réalisation d'une série d'actions, chacune des actions de la série ayant été sélectionnée à l'aide du système de sélection d'action automatisée selon l'une quelconque des revendications 13 à 15 et/ou à l'aide du procédé de sélection d'action automatisée selon la revendication 19 ou 20, dans lequel les paramètres d'état-de-système/contexte (SC(t)) résultant de la réalisation, par ledit robot ou autre dispositif d'auto-développement, d'une action donnée de ladite série ont constitué les paramètres d'état-de-système/contexte (SC(t)) utilisés par le système

IPR2025-00943
Tesla EX1002 Page 1520

ou procédé de sélection d'action automatisée lors de la sélection de l'action suivante de ladite série.

22. Machine de prévision formée en ayant été amenée à prévoir une série d'ensembles de variables prévues (VAR) à l'aide du système de formation de machine de prévision automatisée selon la revendication 11 ou 12, des données indicatives des valeurs réelles des variables prévues (VAR) résultant de la réalisation d'une action donnée, ainsi que des données indicatives des valeurs applicables de tout paramètre d'état-de-système/contexte au moment où ladite action donnée a été effectuée, ayant été retournées à la machine de prévision de sorte qu'une adaptation, répondant aux données de retour, soit constituée de la fonction de prévision appliquée pour les actions candidates de la même région de l'espace d'état-de-système/contexte/action que la région contenant la combinaison de l'action qui a produit les données de retour et tout paramètre d'état-de-système/contexte applicable au moment où ladite action donnée a été réalisée.

23. Système informatique conçu pour prévoir la valeur d'un ensemble d'une ou de plusieurs variables prévues (VAR), ledit système informatique ayant été formé en ayant été amené à prévoir une série d'ensembles de variables prévues (VAR) à l'aide du système de formation de machine de prévision automatisée selon la revendication 11 ou 12, des données indicatives des valeurs réelles des variables prévues (VAR) résultant de la réalisation d'une action donnée, ainsi que des données indicatives des valeurs applicables de tout paramètre d'état-de-système/contexte au moment où ladite action donnée a été effectuée, ayant été retournées vers le système informatique de sorte qu'une adaptation, répondant aux données de retour, soit constituée de la fonction de prévision appliquée pour les actions candidates de la même région de l'espace d'état-de-système/contexte/action que la région contenant la combinaison de l'action qui a produit les données de retour et tout paramètre d'état-de-système/contexte applicable au moment où ladite action donnée a été réalisée.

24. Système expert conçu pour prévoir la valeur d'un ensemble d'une ou de plusieurs variables prévues (VAR), ledit système expert se composant d'une machine de prévision ayant été formée en ayant été amenée à prévoir une série d'ensembles de variables prévues (VAR) à l'aide du système de formation de machine de prévision automatisée selon la revendication 11 ou 12, des données indicatives des valeurs réelles des variables prévues (VAR) résultant de la réalisation d'une action donnée, ainsi que des données indicatives des valeurs applicables de tout paramètre d'état-de-système/contexte au moment où ladite action donnée a été effectuée, ayant été retournées vers ladite machine de précision de sorte qu'une adaptation, répondant aux données de retour, soit constituée de la fonction de prévision appliquée pour les actions candidates de la même région de l'espace d'état-de-système/contexte/action que la région contenant la combinaison de l'action qui a produit les données de retour et tout paramètre d'état-de-système/contexte applicable au moment où ladite action donnée a été réalisée.

IPR2025-00943
Tesla EX1002 Page 1521

FIG.1

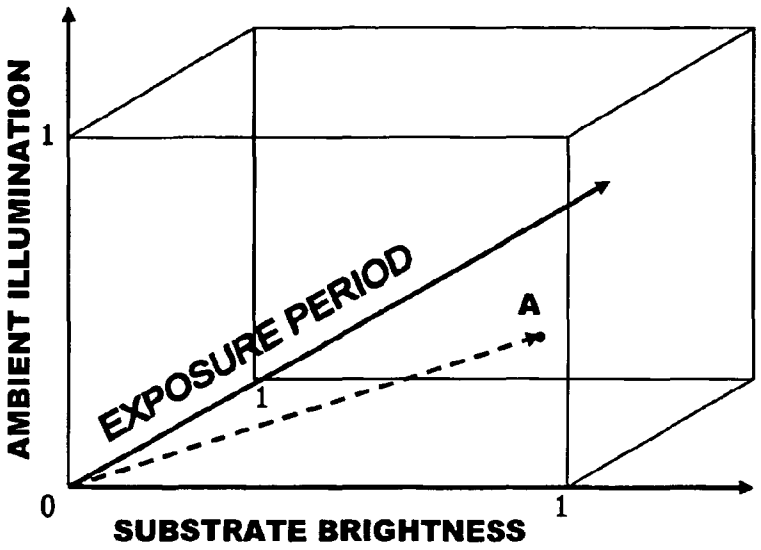
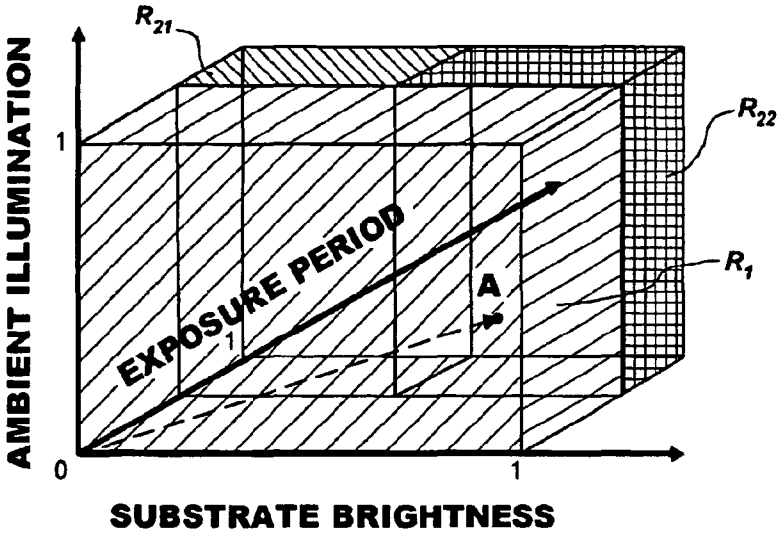


FIG.2



IPR2025-00943
Tesla EX1002 Page 1522

FIG.3

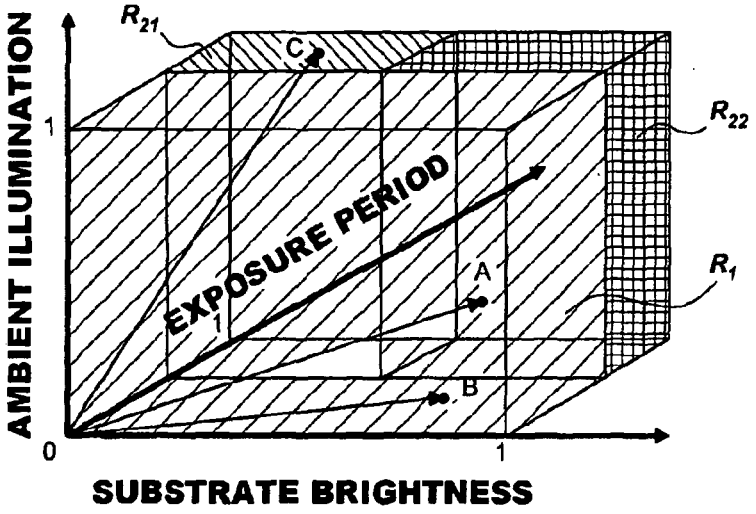
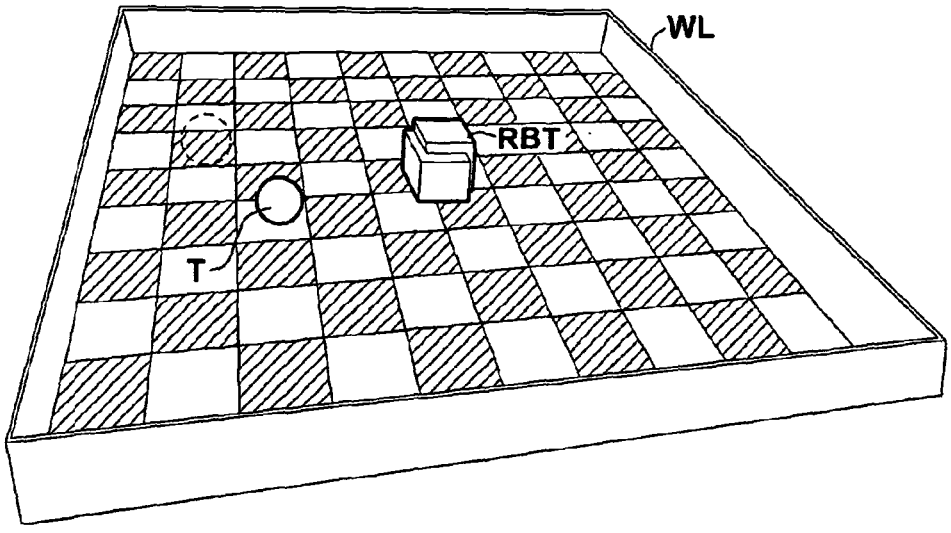
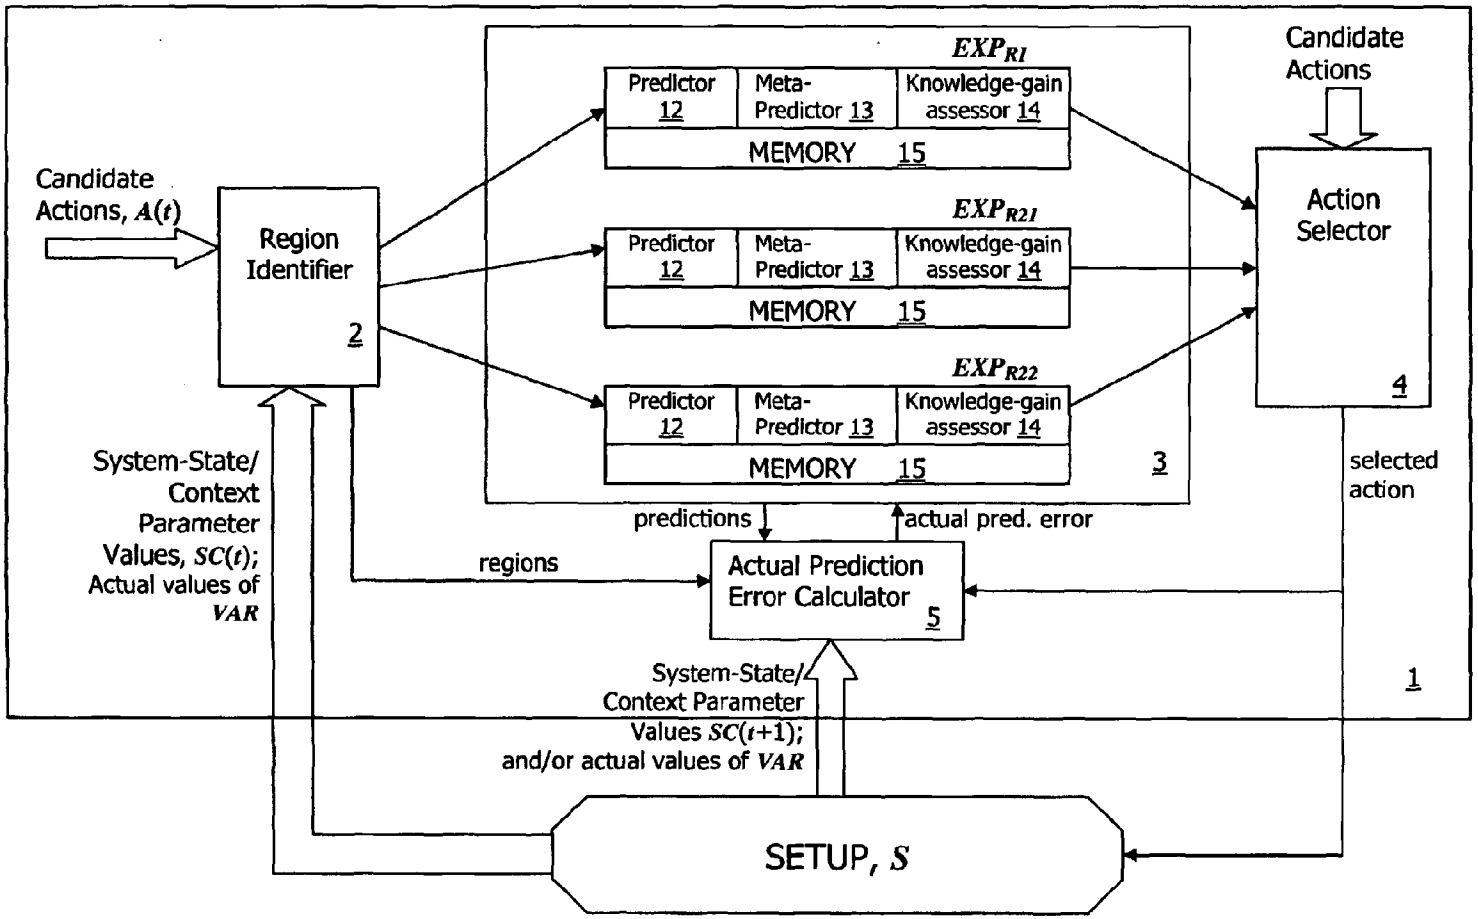


FIG.8



IPR2025-00943
Tesla EX1002 Page 1523

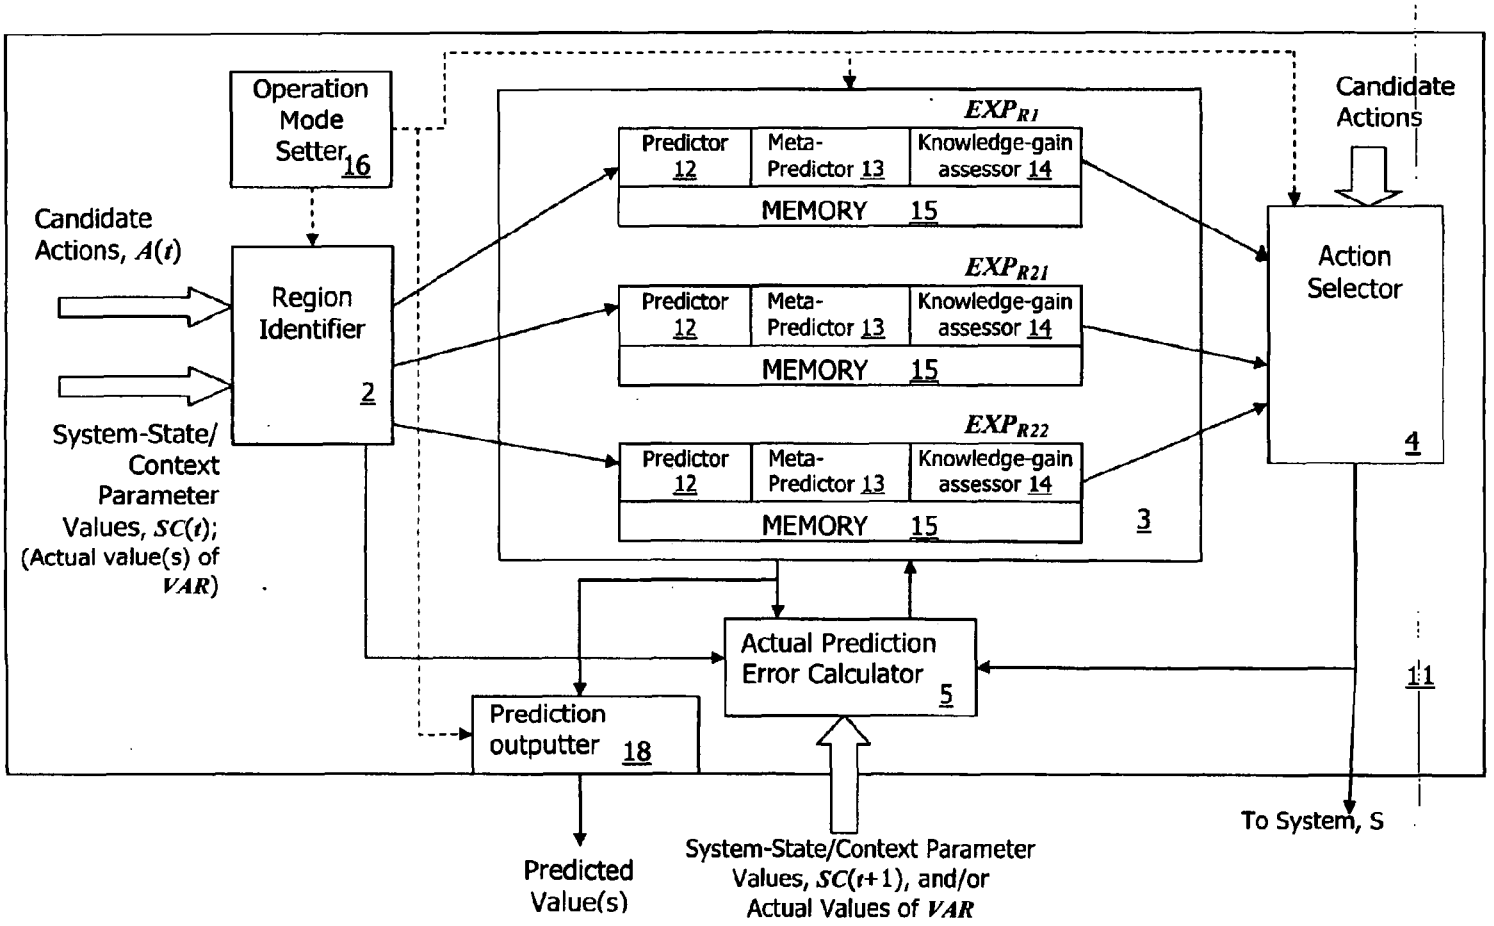
FIG.4



EP 1 622 072 B1

IPR2025-00943
Tesla EX1002 Page 1524

FIG.5



EP 1 622 072 B1

IPR2025-00943
Tesla EX1002 Page 1525

FIG.6

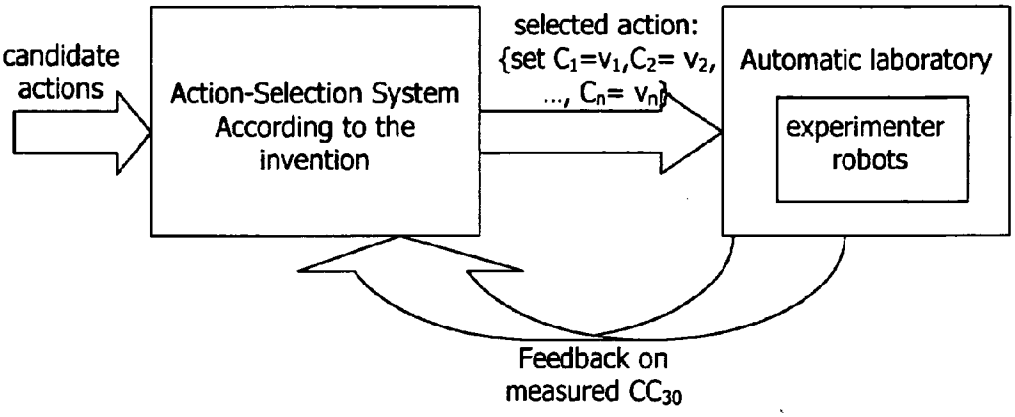
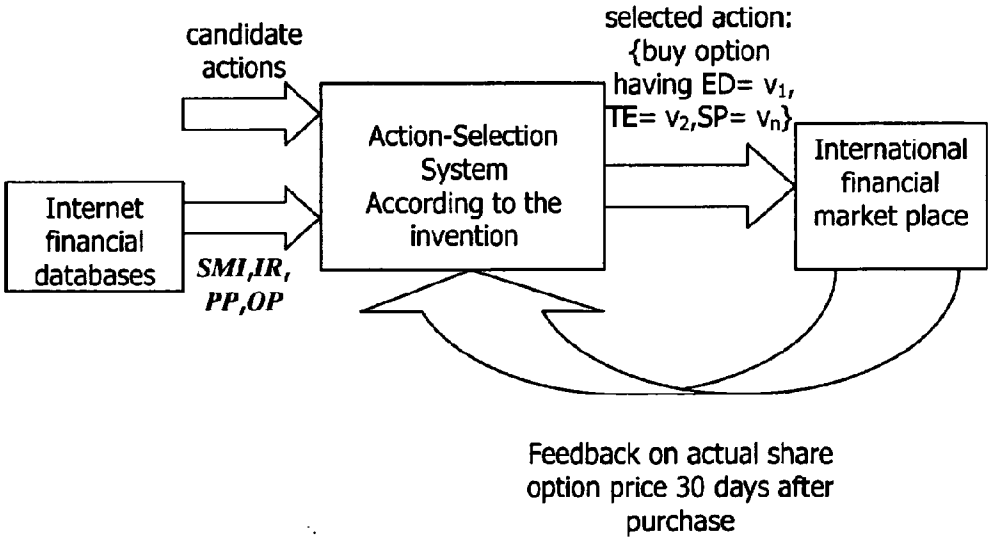
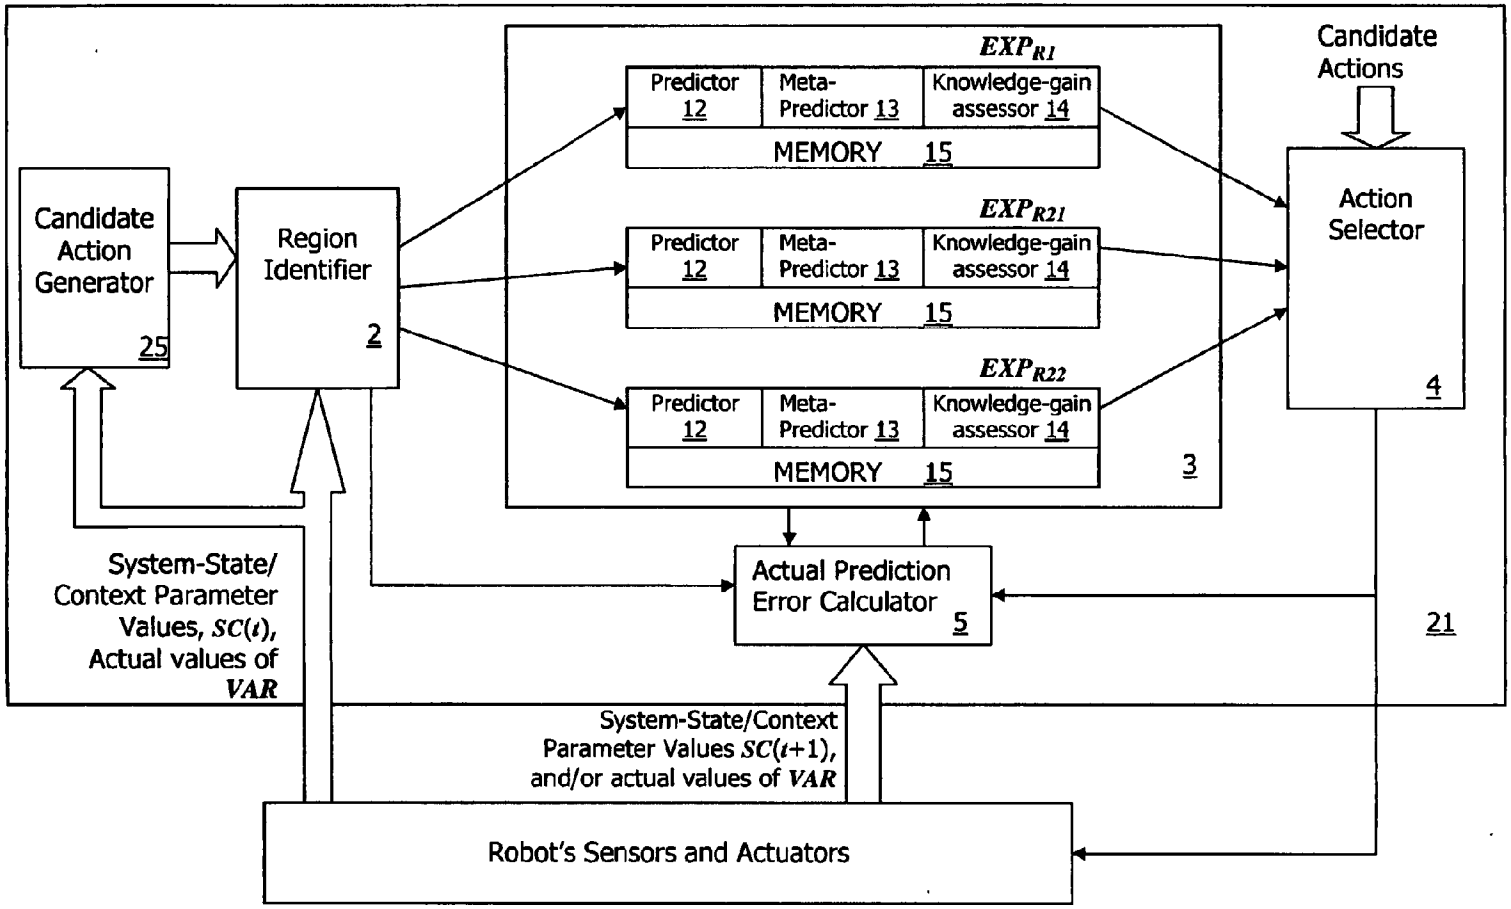


FIG.7



IPR2025-00943
Tesla EX1002 Page 1526

FIG.9



EP 1 622 072 B1

IPR2025-00943
Tesla EX1002 Page 1527

FIG.10

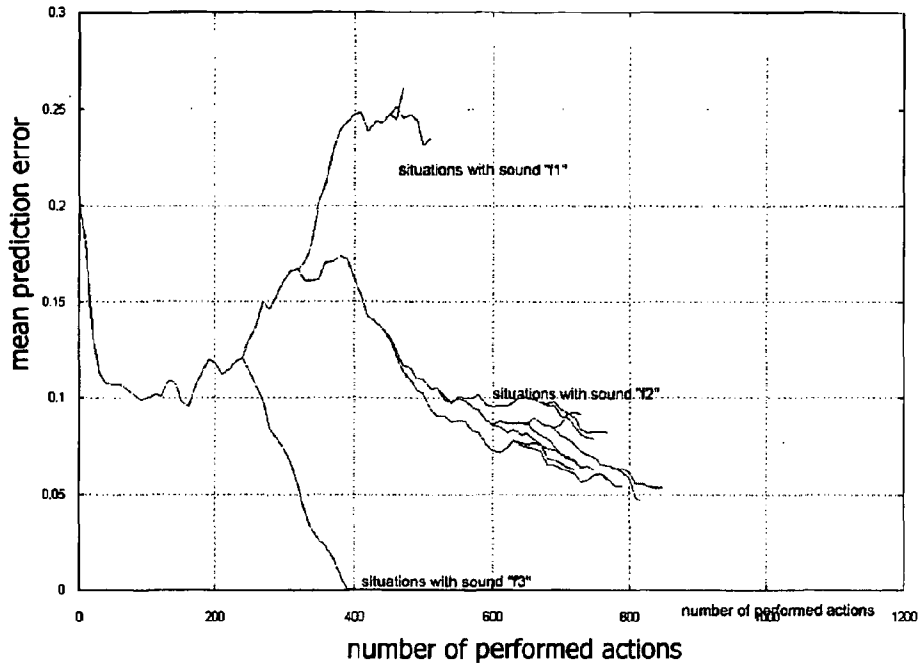
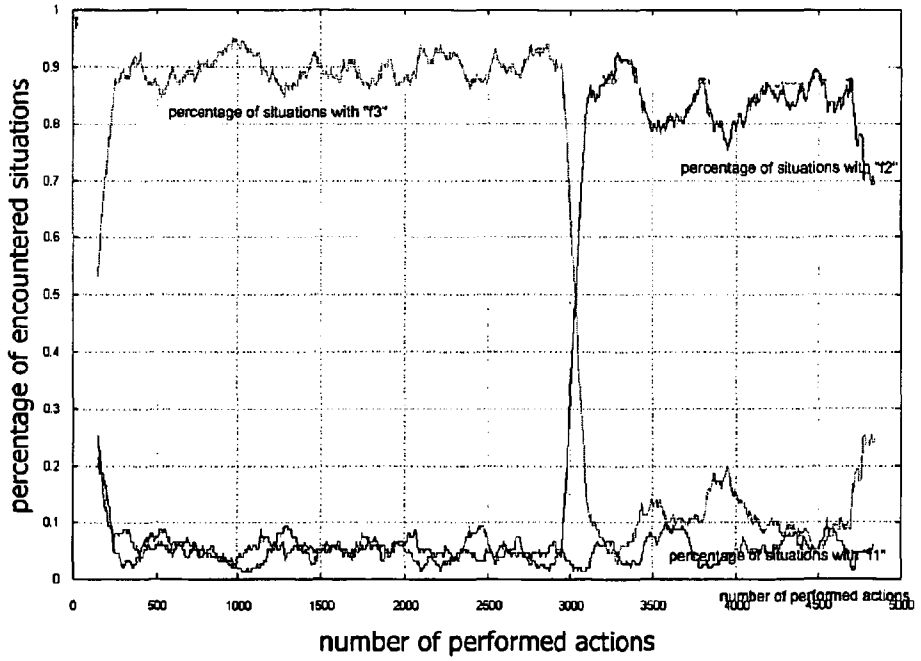
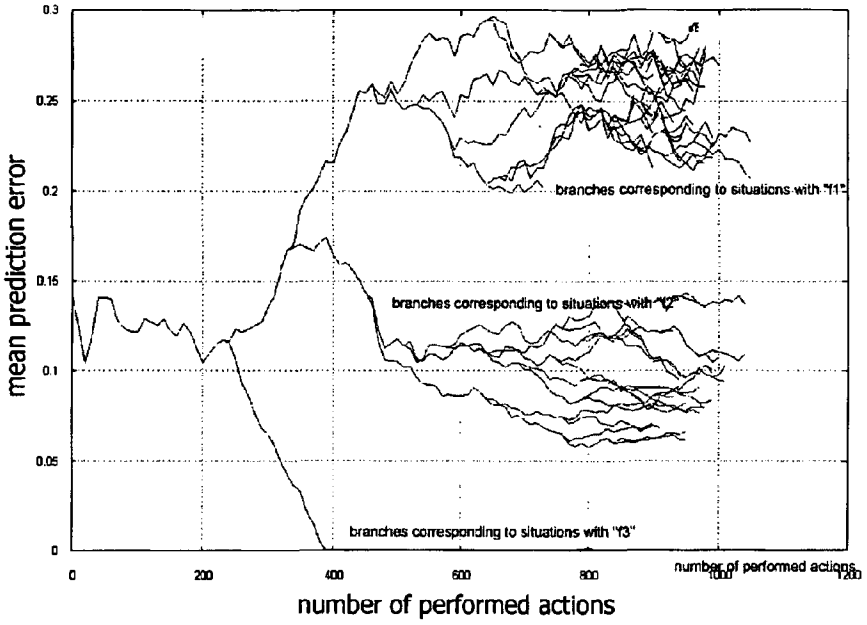


FIG.11



IPR2025-00943
Tesla EX1002 Page 1528

FIG.12



IPR2025-00943
Tesla EX1002 Page 1529

FIG.13

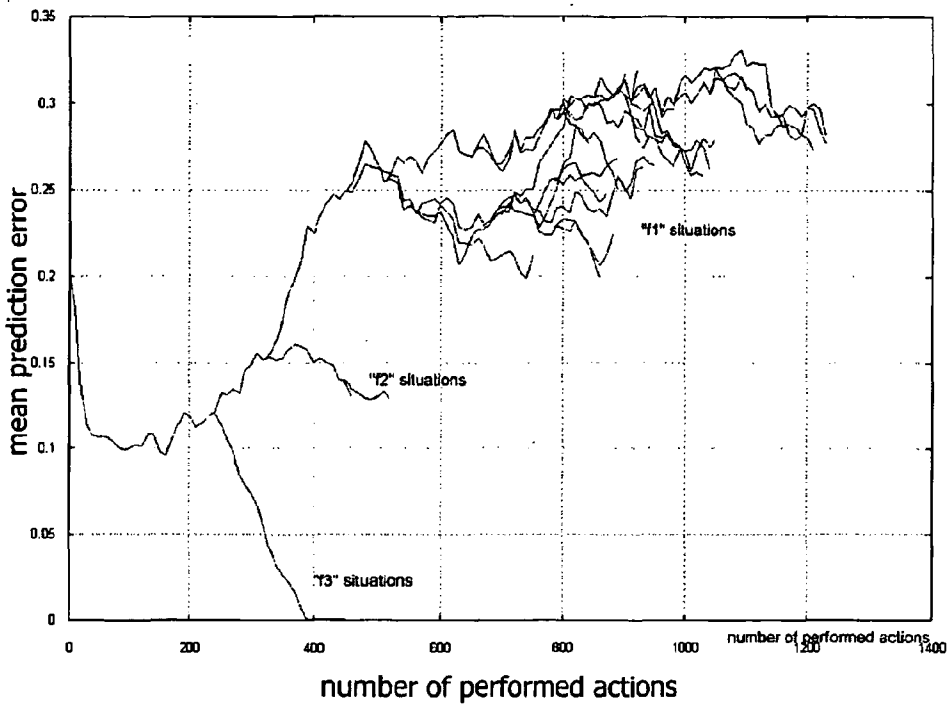
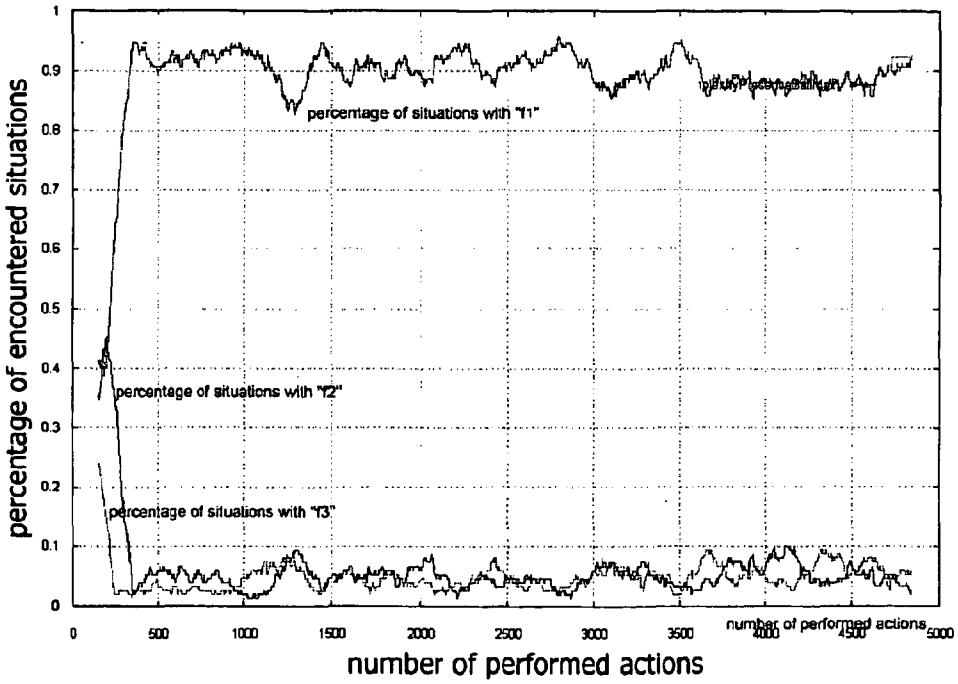


FIG.14



IPR2025-00943
Tesla EX1002 Page 1530

FIG.15

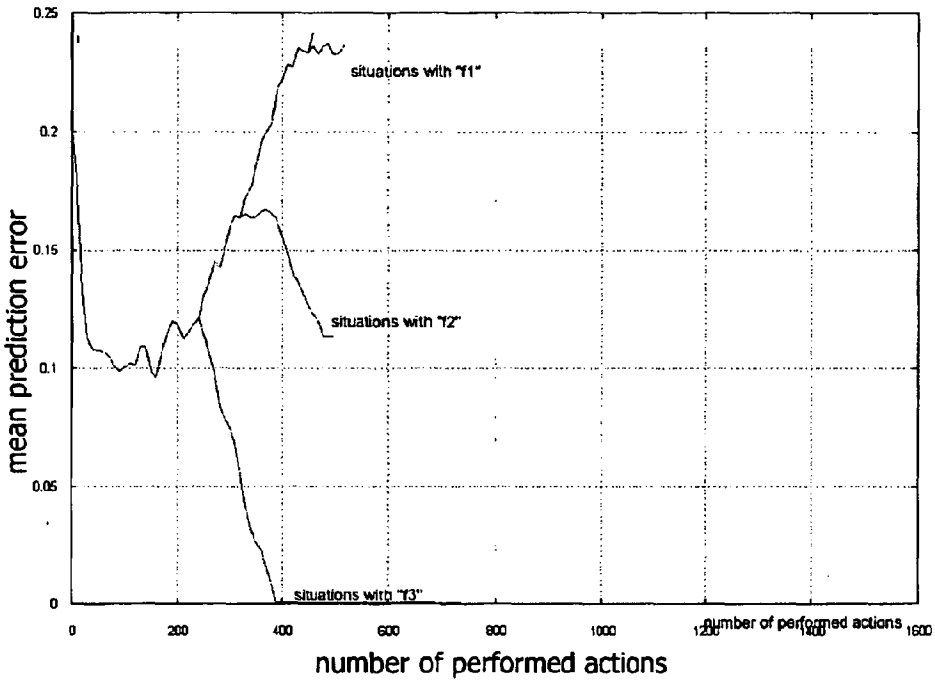
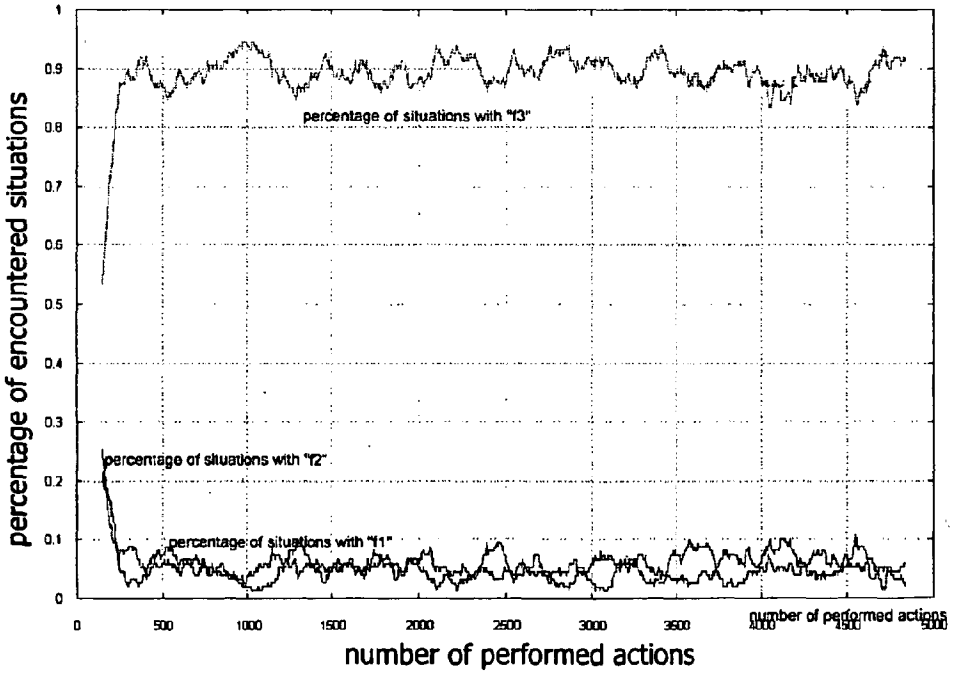
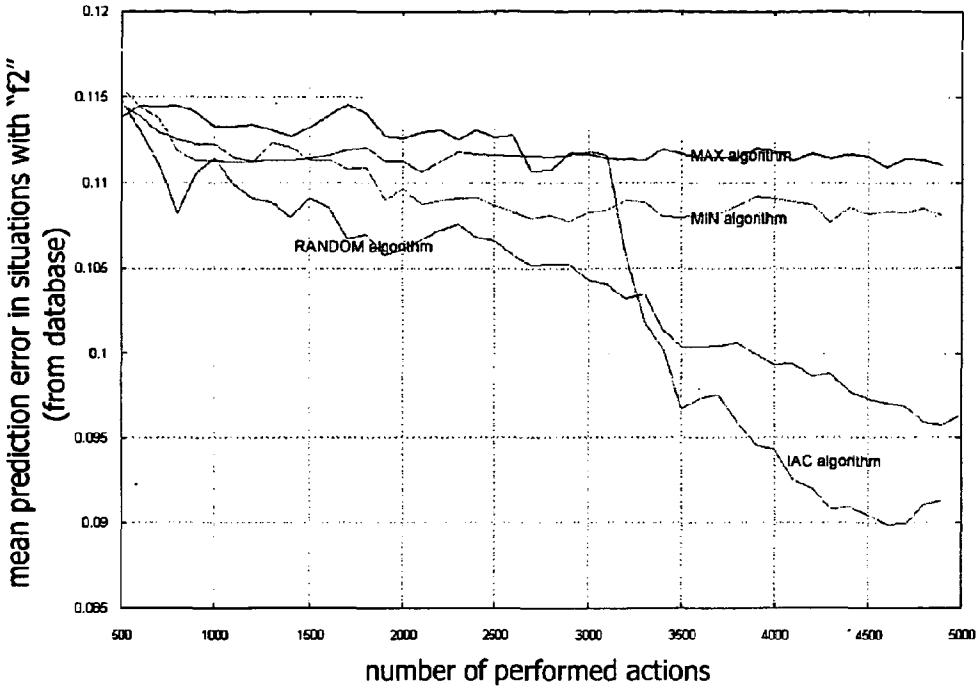


FIG.16



IPR2025-00943
Tesla EX1002 Page 1531

FIG.17



IPR2025-00943
Tesla EX1002 Page 1532

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- EP 03291359 A **[0032]**

Non-patent literature cited in the description

- **J. Denzler ; C. Brown.** Information Theoretic Sensor Data Selection for Active Object Recognition and State Estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 2001, vol. 2, 145-157 **[0022]**

IPR2025-00943
Tesla EX1002 Page 1533

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 727 089 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
29.11.2006 Bulletin 2006/48

(51) Int Cl.:
G06T 7/20 *(2006.01)* ***G05D 1/02*** *(2006.01)*

(21) Application number: **06017695.5**

(22) Date of filing: **27.11.2000**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE TR

(30) Priority: **26.11.1999 US 167587 P**
01.09.2000 US 230166 P

(62) Document number(s) of the earlier application(s) in accordance with Art. 76 EPC:
00980706.6 / 1 257 971

(71) Applicant: **MobilEye Technologies, Ltd.**
1034 Nicosia (CY)

(72) Inventors:
- **Stein, Gideon**
  Jerusalem 96263 (IL)
- **Shashua, Amnon**
  Mevaseret Zion 90805 (IL)
- **Mano, Ofer**
  Modiin 71700 (IL)

(74) Representative: **Skone James, Robert Edmund**
Gill Jennings & Every LLP
Broadgate House
7 Eldon Street
London EC2M 7LH (GB)

Remarks:
This application was filed on 25 - 08 - 2006 as a divisional application to the application mentioned under INID code 62.

(54) System and method for estimating ego-motion of a moving vehicle using successive images recorded along the vehicle's path of motion

(57) An ego-motion determination system is disclosed for generating an estimate as to the ego-motion of a vehicle moving along a roadway. The ego-motion determination system includes an image information receiver and a processor. The image information receiver is configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway. The processor is configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle, including the translation of the vehicle in the forward direction and the rotation of the vehicle around a vertical axis as between, for example, successive images.

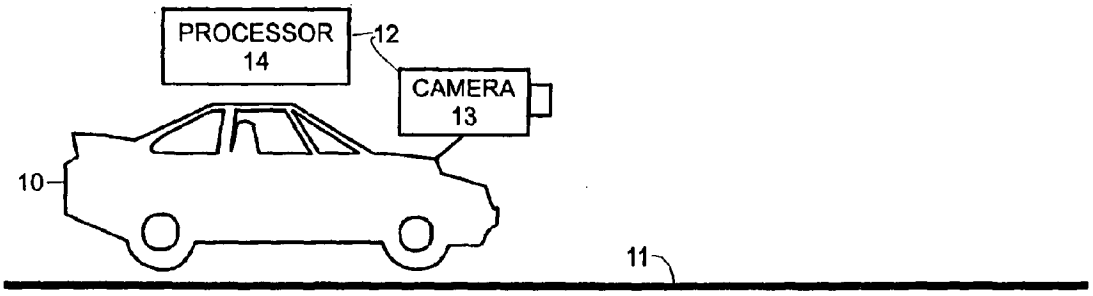


FIG. 1

EP 1 727 089 A2

IPR2025-00943
Tesla EX1002 Page 1534

Description

FIELD OF THE INVENTION

[0001] The invention relates generally to the field of systems and methods for estimating ego-motion (that is, "self-motion) of a moving vehicle, and more specifically to systems and methods that estimate ego-motion using successively-recorded images recorded along the vehicle's path ofmotion.

BACKGROUND OF THE INVENTION

[0002] Accurate estimation of the ego- ("self-") motion of a vehicle relative to a roadway is an important component in autonomous driving and computer vision-based driving assistance. Using computer vision techniques to provide assistance while driving, instead of mechanical sensors, allows for the use of the information that is recorded for use in estimating vehicle movement to also be used in detecting obstacles, identifying lanes and the like, without the need for calibration between sensors as would be necessary with mechanical sensors. This reduces cost and maintenance.

[0003] There are several problems in estimating ego-motion of a vehicle. Typically, roads have few feature points, if any. The most obvious features in a road, such as lane markings, have a generally linear structure, whereas background image structures, such as those associated with other vehicles, buildings, trees, and the like, will typically have many feature points. This will make image- or optical-flow-based estimation difficult in practice. In addition, typically images that are recorded for ego-motion estimation will contain a large amount of "outlier" information that is either not useful in estimating ego-motion, or that may result in poor estimation. For example, in estimating of ego-motion relative to a fixed structure such as a road, images of objects that are moving relative to the road, such as other moving vehicles and even moving wipers, will contribute false information for the ego-motion estimation. In addition, conditions that degrade image quality, such as raindrops and glare, will also make accurate ego-motion estimation difficult.

SUMMARY OF THE INVENTION

[0004] The invention provides a new and improved system and method for estimating ego-motion using successively-recorded images recorded along the vehicle's path of motion..

[0005] In brief summary, the invention provides an ego-motion determination system for generating an estimate as to the ego-motion of a vehicle moving along a roadway. The ego-motion determination system includes an image information receiver and a processor. The image information receiver is configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway. The processor is configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle, including the translation of the vehicle in the forward direction and the rotation of the vehicle around a vertical axis as between, for example, successive images.

[0006] Several methodologies are disclosed for generating the ego-motion estimate of the vehicle.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

FIG. 1 schematically depicts a vehicle moving on a roadway and including an ego-motion estimation system constructed in accordance with the invention;

FIG. 2 is a flow chart depicting operations performed by the ego-motion estimation system in determining ego-motion of the vehicle in accordance with one methodology; and

FIG. 3 is a flow chart depicting operations performed by the ego-motion estimation system in determining ego-motion of the vehicle in accordance with a second methodology.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

[0008] FIG. 1 schematically depicts a vehicle 10 moving on a roadway 11 and including an ego-motion estimation system 12 constructed in accordance with the invention. The vehicle 10 may be any kind of vehicle 10 that may move on the roadway 11, including, but not limited to automobiles, trucks, buses and the like. The ego-motion estimation system 12 includes a camera 13 and a ego-motion estimation system processor 14. The camera 13 is mounted on the vehicle 10 and is preferably pointed in a forward direction, that is, in the direction in which the vehicle would normally

IPR2025-00943
Tesla EX1002 Page 1535

move, to record successive images as the vehicle moves over the roadway. Preferably as the camera 13 records each image, it will provide the image to the ego-motion estimation system processor 14. The ego-motion estimation system processor 14, in turn, will process information that it obtains from the successive images, possibly along with other information, such as information from the vehicle's speedometer (not separately shown) to determine the ego-motion (that is, the self-motion) of the vehicle relative to the roadway 11. The ego-motion estimation system processor 14 may also be mounted in or on the vehicle 11 and may form part thereof. The ego-motion estimates generated by the ego-motion estimation system processor 14 may be used for a number of things, including, but not limited to obstacle and lane detection, autonomous driving by the vehicle, perhaps also using positioning information from, for example, the global positioning system ("GPS") and roadway mapping information from a number of sources known to those skilled in the art, and the like. Operations performed by the ego-motion estimation system processor 14 in determining ego-motion of the vehicle 10 will be described in connection with the flow charts depicted in FIGS. 2 and 3.

[0009] Before proceeding further, it would be helpful to provide background of the operations performed by the ego-motion estimation system processor 14 depicted in FIG. 1. Generally, as between two images Ψ and Ψ' the ego-motion estimation system processor 14 attempts to determine the translation $t = (t_x, t_y, t_z)^T$ and rotation $w = (w_x, w_y, w_z)^T$ (where "T" refers to the transpose operation, t_i refers to translation along the respective "X," "Y" and "Z" axes, and w_i refers to rotation around the respective axis) of the camera 13 affixed to the vehicle 10. Since the camera 13 is affixed to the vehicle 10, the translation and rotation of the camera 13 will also conform to the translation and rotation of the vehicle 10. In that case, for a point $p = (x,y)^T$ (where "x" and "y" are coordinates of a point or feature in the image) that is a projection of a point $P = (X,Y,Z)^T$ (where "X," "Y" and "Z" are coordinates of the point in three-dimensional space), the flow vector for the point, that is, the vector indicating the motion of the same point in three dimensional space from is position in the image Ψ to the image Ψ', has components (u,v)

$$u = \frac{1}{Z} S_1^T t + \left(\hat{p} \, X \, S_1 \right)^T w$$

$$v = \frac{1}{Z} S_2^T t + \left(\hat{p} \, X \, S_2 \right)^T w \qquad (1),$$

where "X" in equations (1) refers to the matrix cross-product and

$$S_1 = \begin{pmatrix} f \\ 0 \\ -x \end{pmatrix} \qquad S_2 = \begin{pmatrix} 0 \\ f \\ -y \end{pmatrix} \qquad \hat{p} = \begin{pmatrix} x/f \\ y/f \\ 1 \end{pmatrix} \qquad (2),$$

where "f" is the focal length of the camera 13, which is presumed to be known. The roadway on which the vehicle 10 is traveling is modeled as a plane. The equation for points on a plane is

$$AX + BY + BZ = 1 \qquad (3),$$

where the "X" and "Y" axes correspond to horizontal "x" and vertical "y" axes of the image plane, and the "Z" axis is perpendicular to the image plane. The camera 13 may be tilted slightly downwardly to increase the amount of the image that corresponds to the road, and reduce the amount of the image that corresponds to other features, such as other traffic, buildings, trees and the like, and in that case, the images Ψ and Ψ' will be rectified so that the images' z axis will be parallel to the plane of the roadway, as will be described below.

[0010] Dividing equation (3) by "Z" provides

IPR2025-00943
Tesla EX1002 Page 1536

$$\frac{1}{Z} = ax + by + c \qquad (4),$$

where $a=A/f$, $b=B/f$ and $c=C$. Substituting equation (4) into equations (1) results in

$$\begin{aligned} u &= (ax+by+c)S_1^T t + \left(\hat{p}\ X\ S_1\right)^T w \\ v &= (ax+by+c)S_2^T t + \left(\hat{p}\ X\ S_2\right)^T w \end{aligned} \qquad (5).$$

Expanding equations (5) results in

$$\begin{aligned} u = & -\left(ct_z + aft_x\right)x + \left(bft_x - w_z\right) + \left(fw_y + cft_x\right) + \\ & \left(\frac{w_y}{f} - at_z\right)x^2 - \left(\frac{w_x}{f} + bt_z\right)xy \end{aligned} \qquad (6)$$

$$\begin{aligned} v = & \left(w_z + aft_x\right)x + \left(-ct_z + bft_y\right)y - \left(fw_x + cft_y\right) + \\ & \left(\frac{w_y}{f} - at_z\right)xy - \left(\frac{w_x}{f} + bt_z\right)y^2 \end{aligned} \qquad (7).$$

Equations (6) and (7) are a special case (the "calibrated camera 13" case) of an eight parameter model for a camera 13 moving relative to a plane:

$$u = \alpha_1 x + \alpha_2 y + \alpha_3 + \alpha_7 x^2 + \alpha_8 xy \qquad (8)$$

$$v = \alpha_4 x + \alpha_5 y + \alpha_6 + \alpha_7 xy + \alpha_8 y^2 \qquad (9)$$

Given the flow vector (u,v), one can recover the parameters α_i, i=1,...,8, from which one can recover the motion parameters t and w.

[0011] One problem arises in connection with the methodology described above in connection with equations (1) through (9), namely, given the large number ofparameters α_i whose values need to be determined, it is difficult to devise a satisfactory method that will reject outliers. This follows from the fact that a relatively large number of optical flow vectors (u,v) will need to be used in determining the values of the parameters, which, in turn, requires a corresponding number of points in both images. In addition, it is difficult to differentiate between optical flow due to rotation around the X and Y axes and translation along the X and Z axes, respectively.

[0012] Accordingly, it is desirable to reduce the number of motion parameters to a minimum. The motion of a vehicle 10 along a roadway can be modeled as being constrained to be a translation along the Z axis, as the vehicle 10 moves

IPR2025-00943
Tesla EX1002 Page 1537

forward or in reverse, and a rotation around the X and Y axes, as the vehicle 10's path deviates from a straight-line course. In that case, equation (5) reduces to

$$u = (ax + by + c)xt_z - \frac{xy}{f}w_x + \left(f + \frac{x^2}{f}\right)w_y$$
$$v = (ax + by + c)yt_z - \left(f + \frac{y^2}{f}\right)w_x + \frac{xy}{f}w_y \qquad \textbf{(10).}$$

If the images are rectified to ensure that the plane of the roadway is parallel to the camera 13's XZ plane, that is, so that the images would be as if the camera 13's optical axis is parallel to the plane of the road, then, in equations (10) a=0 and c=0, in which case

$$u = bxyt_z - \frac{xy}{f}w_x + \left(f + \frac{x^2}{f}\right)w_y$$
$$v = -by^2t_z - \left(f + \frac{y^2}{f}\right)w_x + \frac{xy}{f}w_y \qquad \textbf{(11).}$$

In order to rectify the images, the camera 13 will need to be calibrated. A methodology for calibrating the camera 13 and rectifying the images will be described below.

[0013] In equations (11) there are three motion parameters, t_z (translation along the Z axis), w_x (rotation around the X axis) and w_y (rotation around the Y axis) to be determined from the flow vectors (u,v) associated with points in at least some portions of the images Ψ and Ψ'. Finding corresponding points in the images Ψ and Ψ', that is, points that are projections of the same point in three-dimensional space in the respective images is based on a "photometric constraint"

$$I(x,y,t) - I(x + u\delta t, y + v\delta t, t + \delta t) = 0 \qquad \textbf{(12),}$$

which essentially states that the irradiance, or luminance, of the point $p = (x,y)^T$ in the image Ψ and and the point $p = (x + u\delta t, y + v\delta t)$ in the image Ψ', which are projections ofthe same point in three-dimensional space into the respective images, are the same. In practice, equation (12) will not hold exactly because of noise. If, for every point, the noise is modeled as zero mean Gaussian noise, equation 12 reduces to

$$P\left(I(x,y,t) - I(x + u\delta t, y + v\delta t, t + \delta t)\right) = N(\sigma^2, 0) \qquad \textbf{(13)}$$

and a maximum likelihood can be sought.

[0014] Equation (13) can be computationally intensive, and, instead of using that equation, the motion parameters t_Z, w_X and w_Y can be determined directly from the images by combining the geometric constraints embodied in equation (11) with the photometric constraints embodied in equation (12). In that operation, given two consecutive images Ψ and Ψ', the goal is to determine the probability

IPR2025-00943
Tesla EX1002 Page 1538

$$P\left(\hat{m}|\Psi,\Psi'\right) \quad (14).$$

of that the motion of the vehicle 10 is $\hat{m} = (t_z, w_x, w_y)$ given the two images. The motion m that maximizes (14) is the estimate of the camera 13 motion between the two images Ψ and Ψ', and, thus, the estimate of vehicle 10 motion between the two images.

[0015] According to Bayes' rule,

$$P\left(\hat{m}|\Psi,\Psi'\right) = \frac{P\left(\Psi'|\Psi,\hat{m}\right)P\left(\hat{m}\right)}{P(\Psi')} \quad (15),$$

where $P(\Psi'|\Psi,\hat{m})$ is the probability that, given image Ψ, motion $\hat{m}$ will result in image Ψ', $P(\hat{m})$ is the *a priori* probability that the motion is $\hat{m}$, and $P(\Psi')$ is the *a priori* probability that the image is Ψ'. It will be assumed that $P(\hat{m})$, the probability that the motion is $\hat{m}$, is uniform in a small region $\hat{M}$ around the previous motion estimate, that is, the motion estimate generated as between the "i-1 st" and "i-th" images; it will be appreciated that, if the time period between the times at which images Ψ and Ψ' are recorded is sufficiently small, this assumption will hold. It will further be appreciated that, in equation (15), the denominator $P(\Psi')$ does not depend on the motion $\hat{m}$, and so it does not affect the search for a maximum.

[0016] The probability that, given image Ψ, motion $\hat{m}$ will result in image Ψ', $P(\Psi'|\Psi,\hat{m})$, can be determined by warping the image Ψ' according to the motion $\hat{m}$, thereby to generate a warped image $\hat{\Psi}'$ and determining the sum squared difference ("SSD")

$$S\left(\hat{m}\right) = \frac{1}{N}\sum_{x,y\in R}\left(\hat{\Psi}'(x,y) - \Psi(x,y)\right)^2 \quad (16),$$

between corresponding patches, or regions, R in the images that are believed to be a projection of the roadway in the two images. In equation 16, "N" is the number of points in the region R. It will be appreciated that, if the images Ψ and Ψ' are recorded at times "t" and "t+δt, respectively, the warped image $\hat{\Psi}'$ will represent the image that is assumed would be recorded at time "t" if the motion is $\hat{m}$. Using the SSD criteria (equation 16), $P(\Psi'|\Psi,\hat{m})$, the probability that image Ψ' would result from image Ψ and a motion $\hat{m}$ is given by the probability density function

$$P\left(\Psi'|\Psi,\hat{m}\right) = ce^{-\frac{S\left(\hat{m}\right)}{\sigma^2}} \quad (17),$$

where "c" is a normalization factor and "σ" is the variance of the noise, which is modeled as a zero mean Gaussian function. Since it is assumed that $P(\hat{m})$, the probability that the motion is $\hat{m}$, is uniform in a small region $\hat{M}$ around a previous motion estimate, the problem of finding the maximum likelihood motion $\hat{m}$ for a patch of image Ψ reduces to finding the maximum of the probability density function

$$P\left(\hat{m}|\Psi,\Psi'\right) = ce^{-\frac{S\left(\hat{m}\right)}{\sigma^2}} \quad (18),$$

for motion $\hat{m} \in \hat{M}$.

[0017] Since the motion for which the ego-motion estimation system processor 14 is to generate an estimate is the

IPR2025-00943
Tesla EX1002 Page 1539

translational and rotational motion ofthe vehicle 10 relative to the road, it is desirable for the ego-motion estimation system processor 14 to consider only regions of the images Ψ and Ψ' that comprise projections of the road, and ignore other regions of the images. However, it should be noted that the set R of regions, or patches, of the images that projections of the roadway in the two images Ψ and Ψ', is not known. To accommodate that, instead of attempting to maximize the function defined in equation (18), the image can be tessellated into a set of patches W_i, and a probability density $P(m|W_i,W_i')$ generated for each patch using equations (16) and (18) for the respective patches. The probability density over the entire pair of images Ψ and Ψ' will be

$$P\left(\hat{m}|\Psi,\Psi'\right) = c\frac{\sum_i P\left(\hat{m}|W_i,W_i'\right)\lambda_i\,\alpha_i\,\beta_i}{\sum_i \lambda_i} \qquad (19),$$

where λ_i and α_i are weighting functions whose values generally reflect the confidence that the "i-th" patch is a projection of the road. The value of the gradient strength β_i for a patch reflects the degree to which the patch the contains a texture, and thus will more likely to contain useful information for use in determining ego motion of the vehicle. The motion $\hat{m} \in \hat{M}$ for which equation (19) is the maximum will be deemed to correspond to the actual translational and rotational motion of the vehicle 10 as between the locations at which images Ψ and Ψ' were recorded.

[0018] The weighting function λ_i for the respective "i-th" patch is generated using patches W_i and W'_i from respective images Ψ and Ψ'. In determining the values for the weighting functions λ_i, for patches W_i, W'_i that are not of the road, the motion model reflected in equation (11) is not a good fit; instead, a better fit can be obtained using some other motion of the patch. In addition, for planar image artifacts moving on the roadway surface, such as moving shadows, the maximum of equation (18) will occur far away from the initial guess. Accordingly, the value of the weighting function λ_i for the "i-th" patch W_i, W_i' will correspond to the ratio between the best fit using the motion model in a local region ($\hat{M}$) near the initial guess and the best fit using any motion model over large search region "L." Accordingly, if

$$P_1 = \max\left(\exp\left(-\frac{S_i\left(\hat{m}\right)}{\sigma^2}\right)\right) \qquad (20)$$

for all $\hat{m} \in \hat{M}$ is the value for the best fit in a local search region, where $S_i()$ denotes the SSD over all points in the "i-th" patch, and

$$P_2 = \max\left(\exp\left(-\frac{S_i\left(\hat{m}\right)}{\sigma^2}\right)\right) \qquad (21)$$

for all $\hat{m} \in L$ the value for the best fit over all possible image motions, then

$$\lambda_i = \frac{P_1}{P_2} \qquad (22).$$

Generally, generating the value for P_2 (equation 21) can be computationally intensive. To avoid generating P_2 according to equation (21), the value for P_2 for each patch can be estimated by using the SSD as between a patch in the image

IPR2025-00943
Tesla EX1002 Page 1540

Ψ and the correspondingly-positioned patch in the image Ψ', as well as the SSD's as between the patch in the image Ψ and patches translated horizontally and vertically around the correspondingly-positioned patch in the image Ψ', for a selected number of points. That is, if the patch in image Ψ consists of points p(x,y) centered on p(a,b) (that is, points p(a-α, b-β) through p(a+α,b+β) (α and β being integers), with the patch being of dimensions $2\alpha+1$ by $2\beta+1$), P_2 is generated by using the SSD as between the patch of the same size in image Ψ' consisting of points p(x,y) centered on p(a,b), as well as SSD's as between the patch in image Ψ and patches of the same size in image Ψ' that are centered on points p(a-δ, b-δ) through p(a+δ, b+δ), a total of $(2\delta+1)^2$ patches in image Ψ'. Each patch in image Ψ' can be considered as one of the possible image motions. In one embodiment, δ is selected to be seven, in which case there will be two hundred and twenty five patches in Ψ' for which the SSD will be generated in generating the value for P_2.

[0019] Unlike weighting function λ_i, weighting function α_i for the respective "i-th" patch is generated using only patch W_i and W'_i from image Ψ'. Generally, it should be appreciated that, in three-dimensional space, there are three dominant directions for lines, namely, (i) vertical, for buildings, automobiles, and other objects that extend above the roadway surface, (ii) horizontal lines that are generally parallel to the direction of the roadway, and (iii) horizontal lines that are generally perpendicular to the direction of the roadway. In an image Ψ', vertical lines (that is, lines of type (i)) and horizontal lines that are generally perpendicular to the direction of the roadway (that is, lines of type (iii)) will project in to image Ψ' as vertical and horizontal lines, respectively. On the other hand, horizontal lines that are generally parallel to the direction of the roadway (that is, lines of type (ii)) will appear in image Ψ' as lines that pass through a common point, that is, a vanishing point. It will be appreciated that patches W'_i in image Ψ' that are projections of the roadway will predominately contain lines of type (ii) and (iii). On the other hand, patches W' that are projections of obstacles, such as automobiles will predominately contain lines of type (i) and (iii), while patches W' that are projections of, for example, buildings, fences, and the like, will contain lines of type (i) and (ii).

[0020] Accordingly, the value for weighting function α_i for patch W_i will reflect the degree to which it is deemed to contain projections of lines of type (ii) and (iii), and not projections of lines of types (i) and (iii) or types (i) and (ii). Generally, the directions of lines, if any, passing though a patch can be determined in relation to the gradients of the luminance at the various points in the patch W_i. Each point in the patch W'_i whose gradient (I_x,I_y) is above a selected threshold is considered to lie at or near a line, with the direction of the line being perpendicular to the direction of the gradient. Thus, for those points, the direction of the line associated therewith can be determined, as can whether the line is of type (i), (ii) or (iii). Thus, for each patch W'_i in the image Ψ', three sums S^n_i (n=i, ii or iii) are generated each corresponding to the sum of the magnitudes of the gradients of the points in the patch that are associated with lines of the respective type, the magnitude corresponding to $G=(I_x^2+I_y^2)^{0.5}$. A patch W'_i in image Ψ' is deemed to be:

(a) a projection of the roadway if the sums S^{ii}_i and S^{iii}_i are both large and significantly larger than sum S^i_i, since the sums indicate that the patch is associated with a line or lines that are horizontal and in the direction ofthe roadway and/or perpendicular thereto, but not a line or lines that are vertical;

(b) be a projection of an obstacle, that is, an object generally in the path of the vehicle, if sums S^i_i and S^{iii}_i are both large and significantly larger than sum S^{ii}_i, since the sums indicate that the patch is associated with a line or lines that are vertical and/or horizontal and perpendicular to the direction of the roadway, but not a line or lines that are horizontal and in the direction of the roadway; and

(c) a projection of an object to the side of the path of the vehicle if sums S^i_i and S^{ii}_i are both large and significantly larger than sum S^{iii}_i, since the sums indicate that the patch is associated with a line or lines that are vertical and/or horizontal and in the direction of the roadway.

The value of the weighting function α_i is assigned to the patch based on the degree to which the patch W_i is deemed to be a projection of the roadway (case (a) above).

[0021] It will be appreciated that, if, for a patch W'_i, the sum S^{ii}_i is relatively large, indicating that the patch is associated with a line that is horizontal and in the direction of the roadway, but sums S^i_i and S^{iii}_i are relatively small, indicating that the patch is not associated with a line that is vertical or horizontal and perpendicular to the direction of the roadway, it generally cannot be determined from the set of sums generated for the patch whether the patch is a projection of the roadway (case (a) above) or a projection of an object to the side of the path of the vehicle (case (c) above). However, since the patch is not associated with a line that is vertical, it will generally not be deemed to be a projection of an obstacle (case (b) above). In that case, an assessment as to whether the patch is a projection of the roadway (case (a) above) or a projection of an object to the side of the path of the vehicle (case (c) above) by referring to patches adjacent thereto.

[0022] As noted above, the value of the gradient strength β_i for a patch reflects the degree to which the patch the contains a texture, and thus will more likely to contain useful information for use in determining ego motion of the vehicle. The gradient strength β_i corresponds to

IPR2025-00943
Tesla EX1002 Page 1541

$$\beta_i = \left(\sum_{\hat{m} \in L} \exp\left(- \frac{S_i(\hat{m})}{\sigma^2} \right) \right)^{-1} \quad (23).$$

For relatively uniform patches, the value of the SSD $S_i(\hat{m})$ will be relatively low for all motions, in which case the value of β_i will be relatively low. On the other hand, for patches with texture, the value of the SSD will be relatively high for most motions, in which case the value of β_i. will be relatively high.

[0023] With this background, operations performed by the ego-motion estimation system processor 14 will be describe in connection with the flow chart depicted in FIG. 2. In connection with FIG. 2, it is assumed that the ego-motion estimation system processor 14 already has image Ψ, which it may have used in connection with determining the translational and rotational motion up to the location at which image Ψ was recorded. With reference to FIG. 2, after the ego-motion estimation system processor 14 has received image Ψ' (step 100), it will rectify the image according to information provided during the camera 13 calibration operation (described below) to provide that the optical axis is parallel to the plane defined by the roadway (step 101). In addition, the ego-motion estimation system processor 14 will generate an initial guess as to the translational and rotational motion, using the previous motion estimate and, perhaps information from other sensors if available (step 102). For example, the ego-motion estimation system processor 14 may make use of information from the vehicle 10's speedometer, as well as information as to the time period between the time at which image Ψ was recorded and the time at which image Ψ' was recorded, in generating the initial guess. Generally, it will be appreciated that the time period will be fixed, and will preferably the same for each successive pair of images Ψ and Ψ'. After the ego-motion estimation system processor 14 has generated the initial guess, it will use the initial guess to warp image Ψ' toward image Ψ, thereby to generate a warped image $\hat{\Psi}'$ (step 103).

[0024] After the ego-motion estimation system processor 14 has generated the warped image (step 103), it will divide the image Ψ and the warped image $\hat{\Psi}'$ into patches and, for each pair of corresponding patches in the two images Ψ and $\hat{\Psi}'$, generate the weighting value λ_i. In that operation, the ego-motion estimation system processor 14 will select a patch in the image Ψ (step 104) and generate values for P_2 (step 105), P_1 (equation 20) (step 106) and λ_i (equation 22) (step 107) as described above. In addition, the ego-motion estimation system processor 14 can generate the value for β_i (equation 23) and α_i (step 108). After the ego-motion estimation system processor 14 has generated performed steps 105 through 108 for the selected patch, it will determine whether all of the patches in image Ψ have been processed (step 109) and if not, return to step 104 to select another patch and perform steps 105 through 109 in connection therewith.

[0025] The ego-motion estimation system processor 14 will perform steps 104 through 109 in connection with each patch in the image Ψ. After the ego-motion estimation system processor 14 has performed steps 104 through 109 in connection with all of the patches in the image Ψ, it will sequence from step 109 to step 110 to search for the motion $\hat{m}$ that maximizes the value provided by equation (19) (step 110). That motion $\hat{m}$ will comprise values for translation t_z and rotation w_x, w_Y parameters that will constitute the estimate of the motion of the vehicle 10 as between the point in time at which image Ψ was recorded and the point in time at which image Ψ' is recorded. The ego-motion estimation system processor 14 can perform operations described above in connection with each successive pair of images Ψ and Ψ' to estimate the motion of the vehicle 10.

[0026] In performing steps 106 (to generate the values for P_1) and 110 (to determine the motion $\hat{m}$ that maximizes the value provided by equation (19)), the ego-motion estimation system processor 14 can perform a gradient descent that is limited to a selected cube-shaped region around the initial guess.

[0027] In determining the initial guess (step 102) for each new image Ψ', the ego-motion estimation system processor 14 can use the estimate of the motion generated for the previously-received image. In addition, the size of the region $\hat{M}$ can be adjusted adaptively.

[0028] A second methodology for determining the ego-motion of a vehicle 10 will be described in connection with FIG. 3. As with the methodology described above in connection with FIG. 2, it is assumed that the images have been rectified so that the image planes are perpendicular to the plane represented by the roadway on which the vehicle 10 is traveling, and their horizontal ("x") axes are parallel to the plane represented by the roadway. By way of background, in that case, the equation of a plane

$$AX + BY + CZ = 1 \quad (24)$$

(reference equation (3)) reduces to

IPR2025-00943
Tesla EX1002 Page 1542

$$BY = 1 \qquad (25),$$

[0029] 1 in which case $Y = \frac{1}{|B|}$ is the height of the optical axis of the camera 13 (or, more specifically, the Z axis) above the road. Since, for a point *p(x,y)* in an image that is a projection of a point *P(X,Y,Z)* in three-dimensional space, $y = f\frac{Y}{Z}$, equation (25) becomes

$$\frac{1}{Z} = by \qquad (26)$$

The brightness constraint is

$$uI_x + vI_y + I_t = 0 \qquad (27)$$

for each point, where, at each point (x,y) in the image, I_x and I_y are the horizontal and vertical components of the spatial gradient of the luminance and I_t is the time gradient of the luminance. In addition, the equations for the components *(u, v)* of the flow vector (reference equation (1)) can be written

$$\begin{aligned} u &= \frac{1}{Z}S_1^T t + S_1^T [w]_X p \\ v &= \frac{1}{Z}S_2^T t + S_2^T [w]_X p \end{aligned} \qquad (28)$$

where $[w]_x$ is a skew-symmetric matrix, in which case

$$\frac{1}{Z}S^T t + S^T [w]_X p + I_t = 0 \qquad (29),$$

where

$$S = \begin{pmatrix} fI_x \\ fI_y \\ -xI_x - yI_y \end{pmatrix} \qquad (30).$$

For motion constrained to a plane, equation (29) reduces to

IPR2025-00943
Tesla EX1002 Page 1543

$$(ax + by + c) S^T t + S^T [w]_X p + I_t = 0 \quad (31).$$

Since the images are rectified, equation (26) will hold. In addition, since rotational motion is constrained so as to occur only around the vertical ("Y") axis, $w_x=w_z=0$. Accordingly, equation (31) will correspond to

$$byS^T t + S^T \begin{pmatrix} 0 & 0 & -w_Y \\ 0 & 0 & 0 \\ -w_Y & 0 & 0 \end{pmatrix} \begin{pmatrix} x/f \\ y/f \\ 1 \end{pmatrix} + I_t = 0 \quad (32).$$

Expanding equation 32, and using equation (30),

$$byS^T t - (fI_x + xI_z + yI_y) w_Y + I_t = 0 \quad (33)$$

Generally, the portion u_{wY} of the vertical component of the flow vector that arises due to rotation (w_Y) of the vehicle 10 will be larger than the portion u_{tX} that arises due to translation (t_X) along the "X" axis, and the portions will have a very different form. The portion u_{tX} that arises due to translation along the "X" axis will depend on the distance that the vehicle 10 moves between the times the images are recorded, and will be larger for points towards the bottom of the image than for points further up. In that case, expressions for u_{wY} and u_{iX} are

$$u_{wY} = \left(\frac{x^2}{f} + f \right) w_Y \approx f w_Y \quad (34)$$

$$u_{tX} = \frac{1}{Z} f t_X = byf t_X \quad (35)$$

Since, for small rotations around the vertical "Y," axis, the angle of rotation w_Y is approximately

$$w_Y \approx \frac{t_X}{t_Z} \quad (36),$$

$$\frac{u_{wY}}{u_{tX}} = \frac{f \frac{t_X}{t_Z}}{\frac{1}{Z} f t_X} = \frac{Z}{t_Z} \quad (37)$$

It will be appreciated that values for t_Y, the component of the translation t in the vertical direction, and w_X and w_Z, the X and Z components of rotation w, will be zero. Accordingly, after the ego-motion estimation system processor 14 receives

IPR2025-00943
Tesla EX1002 Page 1544

a new image Ψ', it will determine the values for t, and t_3, the components of the translation t in the forward (along the Z axis) and side (along the X axis) directions, and w_Y, the component of rotation around the vertical (Y) axis. In that operation, the ego-motion estimation system processor 14 will generate an initial estimate as to the motion (step 150, FIG. 3) and use that initial estimate to generate a warped image $\hat{\Psi}'$ (step 151). The ego-motion estimation system processor 14 can use information from a number of sources in connection with generating the initial estimate (step 150), including information from, for example, the vehicle 10's speedometer. Thereafter, the ego-motion estimation system processor 14 divides the image Ψ and $\hat{\Psi}'$ into patches (step 152) and determines which patches are likely to comprise images of the roadway (step 153). In determining which patches are likely to comprise images of the roadway (reference step 153), the ego-motion estimation system processor 14 can generate an SSD (equation 16) for corresponding patches in the images Ψ and $\hat{\Psi}'$ and determine that patches in the warped image $\hat{\Psi}'$ that comprise images of the roadway will be those patches with a relatively high SSD value. The ego-motion estimation system processor 14 then uses the patches identified in step 153 to minimize a cost function of the form

$$\min_{t} \sum_{x,y \in R} |At - b|^2 + \lambda_1 \sum_i \left| W_1 \left(\frac{dt_i}{dT} \right) \right|^p + \lambda_2 \left| W_2 \left(t - t_0 \right) \right|^p \qquad (38),$$

where W_1 and W_2 are weighting matrices that essentially describe confidence in the priors or smoothness values (step 154). Equation (38) can be formalized in the form of a Kalman filter, and the value of "p" can be selected to be one or two depending on whether the L_1 or L_2 norm is to be used.

[0030] As noted above, for methodologies described above in connection with both FIGS. 2 and 3, the ego-motion estimation system processor 14 will initially rectify the images as received from the camera 13. In the above description, the images I and I' are images as rectified by the ego-motion estimation system processor 14. To rectify the images as received from the camera 13, the camera 13 will need to be calibrated during a calibration operation prior to use in connection with recording images for use in estimating vehicle 10 motion as described above. Before describing operations to be performed during calibration, it would be helpful to consider the effects of incorrect calibration. If, for example, a vehicle is driving along a straight road, and if the optical axis of the camera is aligned with the direction of motion, the flow field of successively recorded images will be an expansion field with the focus of expansion located at the center of the respective images, that is, at the origin (x,y)=(0,0) of the image's rectilinear coordinate system. On the other hand, if the camera is mounted on the vehicle with a small rotation around the vertical ("Y") axis in three-dimensional space, then the focus of expansion will be displaced along the image's horizontal ("x") axis. In that case, the motion model defined by equation (11) will not account for the flow field, but will be well approximated by a forward translation and a rotational velocity w_y around the vertical ("Y") axis.

[0031] Accordingly, errors in the orientation of the camera around the vertical ("Y") axis in three dimensional space will create a bias in the rotation estimate, in which case a curved path would be estimated instead of a straight path. Similarly, errors in the camera's orientation around the horizontal ("X") axis in three dimensional space will cause a bias in the pitch estimate. Based on these observations, a calibration operation can be performed by having the camera record a sequence of images while the vehicle is being driven down a straight roadway. The vehicle's ego-motion is estimated as described above in connection with FIGS. 2 or 3, and calibration parameters are estimated that would cause the ego-motion to integrate into a straight path.

[0032] The invention provides a number of advantages. In particular, the invention provides an arrangement for determining ego-motion of a vehicle 10 on a roadway from a series of images recorded by a camera 13 mounted on the vehicle 10, at least a portion of the images comprising projections of the roadway, and without requiring mechanical sensors which are normally not provided with a vehicle 10 and that would, if provided, increase the cost and maintenance expenses thereof.

[0033] It will be appreciated that a system in accordance with the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof, any portion of which may be controlled by a suitable program. Any program may in whole or in part comprise part of or be stored on the system in a conventional manner, or it may in whole or in part be provided in to the system over a network or other mechanism for transferring information in a conventional manner. In addition, it will be appreciated that the system may be operated and/or otherwise controlled by means of information provided by an operator using operator input elements (not shown) which may be connected directly to the system or which may transfer the information to the system over a network or other mechanism for transferring information in a conventional manner.

[0034] The foregoing description has been limited to a specific embodiment of this invention. It will be apparent, however, that various variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention. It is the object of the appended claims to cover these and such other variations and

IPR2025-00943
Tesla EX1002 Page 1545

modifications as come within the true spirit and scope of the invention.

Claims

1. A method of controlling a vehicle moving along a roadway, the method comprising:

   acquiring first and second camera images of the roadway at different times;
   providing a model that models the ego-motion of the vehicle between the times that the first and second images are acquired that is a function of only three motion parameters;
   determining values for the three motion parameters that maximize a probability function, which indicates for a given set of values for the motion parameters how probable the first image warps into the second image; and
   controlling the vehicle responsive to the determined variables.

2. A method according to claim 1 wherein controlling the vehicle comprises using an autonomous driving controller that controls the vehicle responsive to the variables.

3. A method according to claim 1 wherein controlling the vehicle comprises controlling the vehicle responsive to a computer vision-based driver assistance that generates driver assistance cues responsive to the determined variables.

4. A method according to any of claims 1-3 wherein the motion parameters comprise a translational velocity parallel to the plane of the roadway and collinear with a direction of motion of the vehicle on a straight roadway, an angular velocity perpendicular to the translational velocity and parallel to the plane of the roadway and an angular velocity perpendicular to the translational velocity and the roadway plane.

5. A method according to any of claims 1-3 wherein the motion parameters comprise a translational velocity parallel to the plane of the roadway and collinear with a direction of motion of the vehicle, on a straight roadway, a second translational velocity parallel to the plane of the roadway and perpendicular to the first translational velocity and an angular velocity perpendicular to the plane of the roadway.

6. A method according to any of the preceding claims wherein the camera has an optic axis and comprising rectifying the images so that the rectified images are acquired as it the optic axis of the camera is parallel to the translational velocity collinear the direction of motion of the vehicle on a straight roadway.

7. A method according to any of the preceding claims wherein the probability function is responsive to a plurality of probability functions, each of which indicates for a given set of values for the motion parameters, how probable a different region of the first image warps into a corresponding region in the second image; and

8. A method according to claim 7 and comprising tessellating the images to determine the different corresponding regions.

9. A method according to claim 7 or claim 8 and comprising weighting the probability functions of the different corresponding regions.

10. A method according to claim 9 wherein weighting comprises weighting responsive to features in the regions.

11. A method according to claim 9 or claim 10 wherein weighting comprises weighting responsive to texture of the regions.

12. A method according to any of claims 9-11 wherein weighting comprises weighting the probability functions for corresponding regions in the first and second images responsive to a measure as to how well the region in the second image warps to its corresponding region in the first image for an ego-motion dependent on only three motion parameters.

13. Apparatus for controlling a vehicle moving along a roadway comprising:

   a camera (13) that acquires images of the roadway at different times;
   a processor (14) that receives images acquired by the camera and:

IPR2025-00943
Tesla EX1002 Page 1546

models the ego-motion of the vehicle between the times that the images are acquired as a function of only three motion parameters;
determines values for the three motion parameters that maximize a probability function, which indicates for a given set of values for the motion parameters how probable the first image warps into the second image; and
controls the vehicle responsive to the determined variables and/or generates driver assistance cues to assist the vehicles driver.

IPR2025-00943
Tesla EX1002 Page 1547

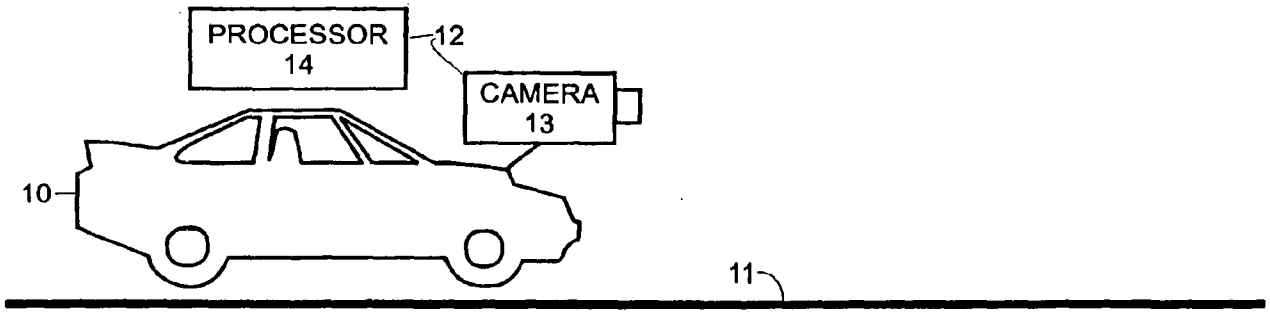
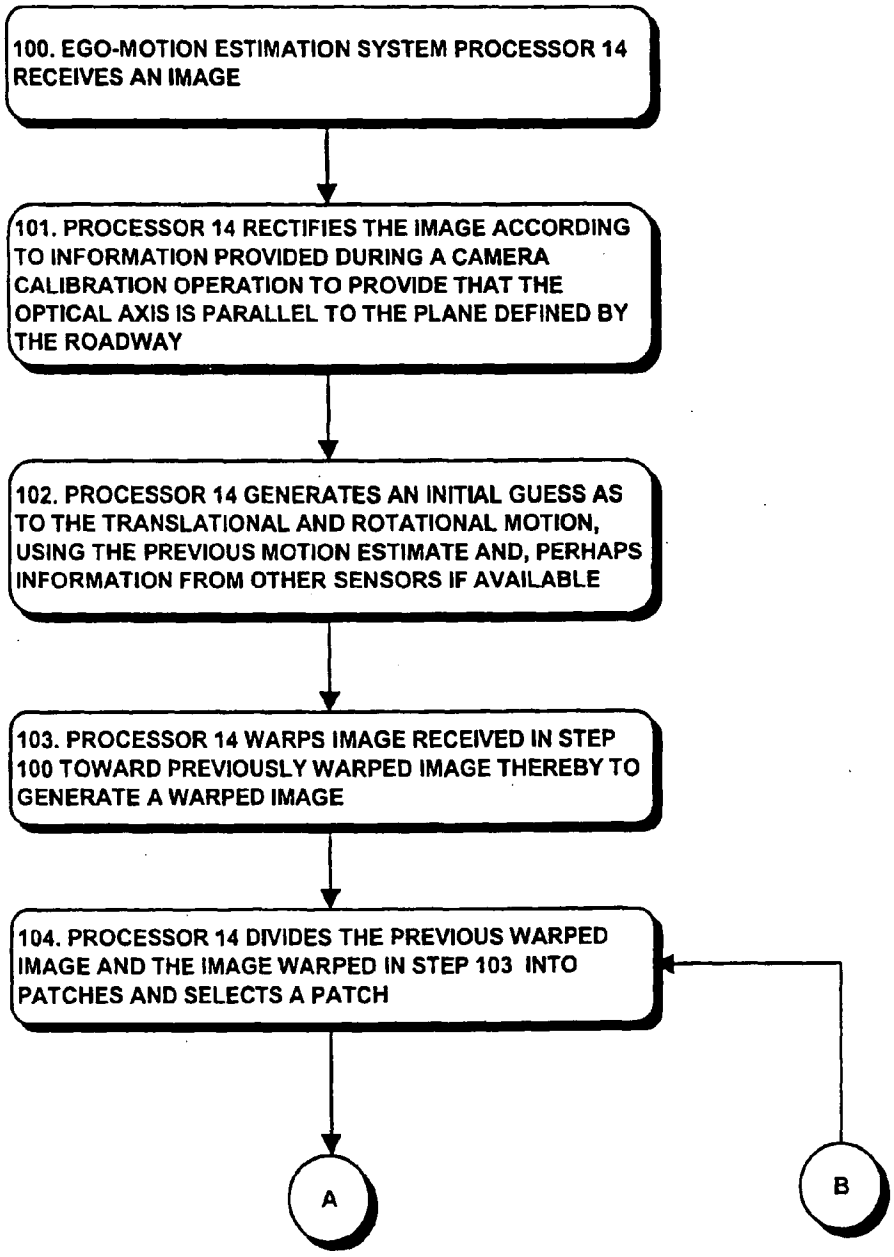


FIG. 1

IPR2025-00943
Tesla EX1002 Page 1548

FIG. 2



IPR2025-00943
Tesla EX1002 Page 1549

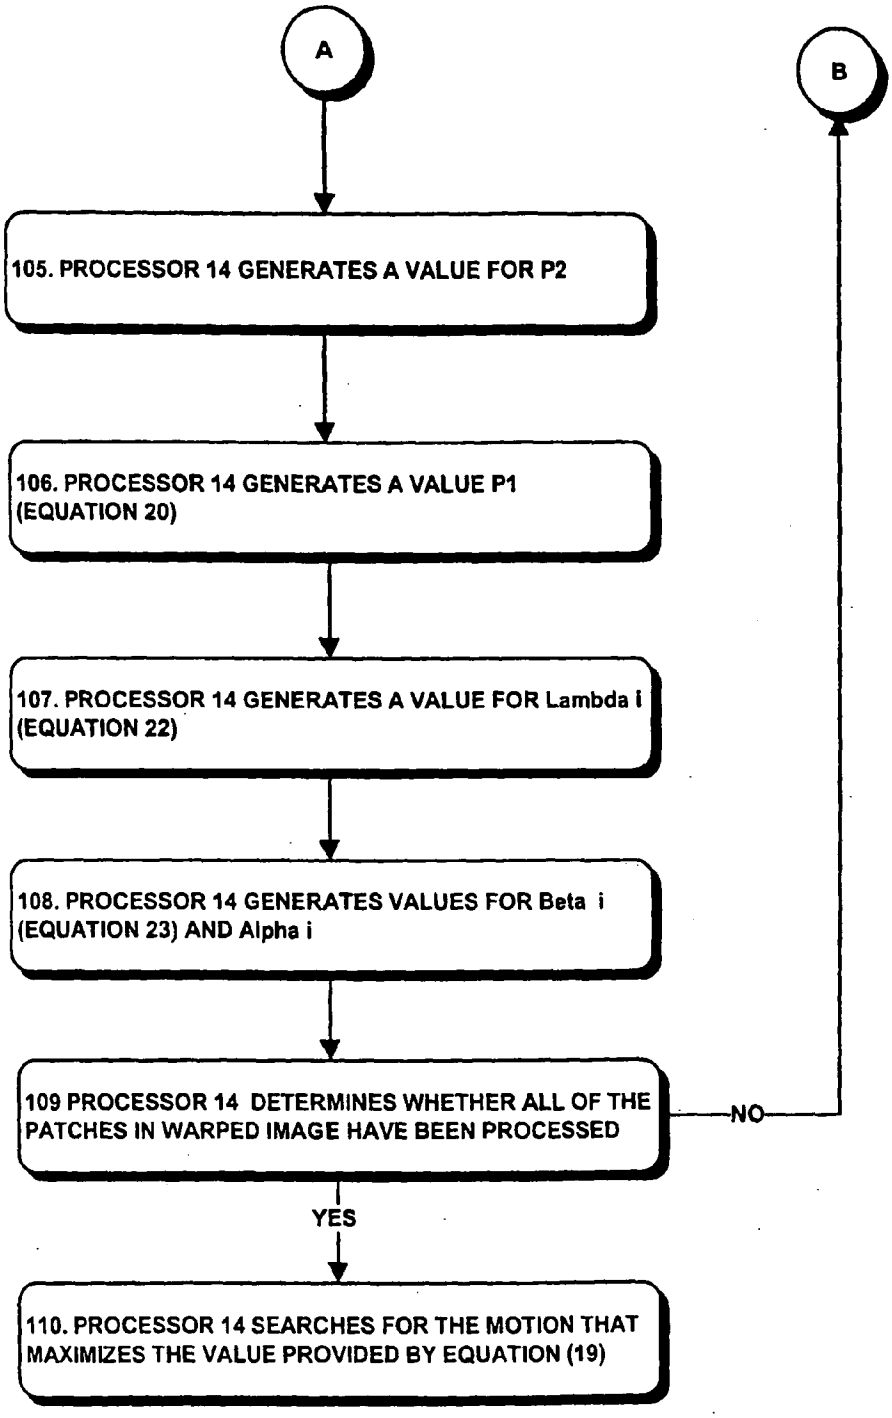
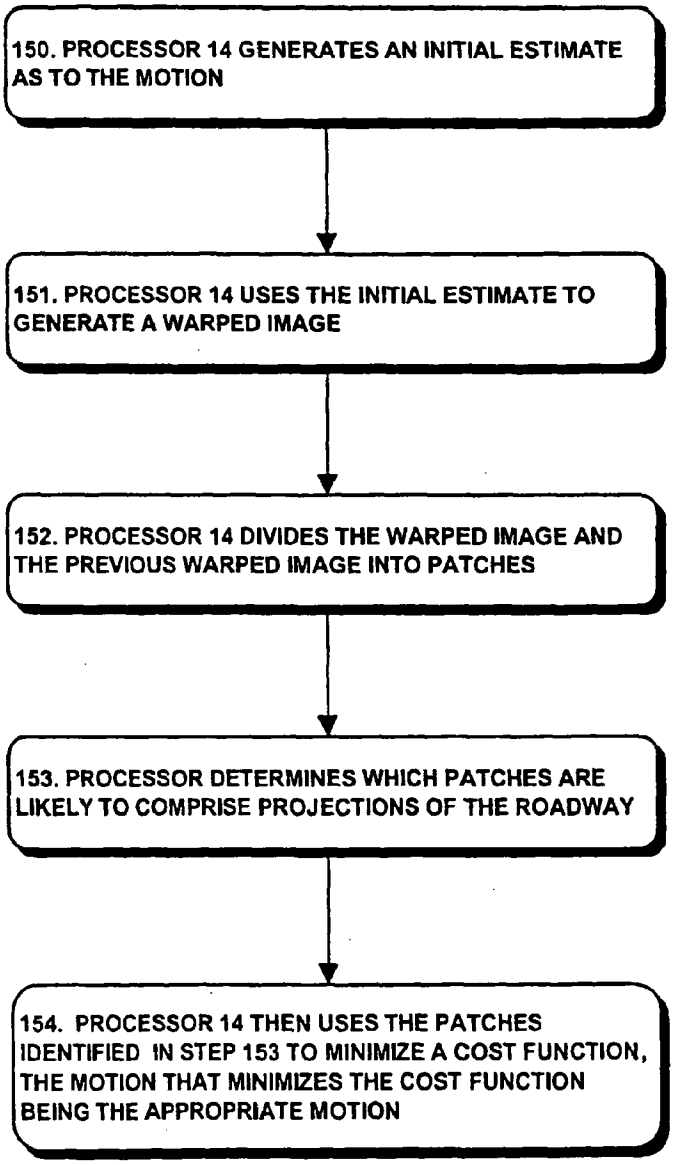


FIG. 2A

IPR2025-00943
Tesla EX1002 Page 1550

FIG. 3



IPR2025-00943
Tesla EX1002 Page 1551

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 727 089 A3

(12) EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
19.09.2007 Bulletin 2007/38

(43) Date of publication A2:
29.11.2006 Bulletin 2006/48

(21) Application number: **06017695.5**

(22) Date of filing: **27.11.2000**

(51) Int Cl.:
G06T 7/20 *(2006.01)* ***G05D 1/02*** *(2006.01)*

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE TR

(30) Priority: **26.11.1999 US 167587 P**
01.09.2000 US 230166 P

(62) Document number(s) of the earlier application(s) in accordance with Art. 76 EPC:
00980706.6 / 1 257 971

(71) Applicant: **MobilEye Technologies, Ltd.**
1034 Nicosia (CY)

(72) Inventors:
- **Stein, Gideon**
  Jerusalem 96263 (IL)
- **Shashua, Amnon**
  Mevaseret Zion 90805 (IL)
- **Mano, Ofer**
  Modiin 71700 (IL)

(74) Representative: **Skone James, Robert Edmund**
Gill Jennings & Every LLP
Broadgate House
7 Eldon Street
London EC2M 7LH (GB)

(54) System and method for estimating ego-motion of a moving vehicle using successive images recorded along the vehicle's path of motion

(57) An ego-motion determination system is disclosed for generating an estimate as to the ego-motion of a vehicle moving along a roadway. The ego-motion determination system includes an image information receiver and a processor. The image information receiver is configured to receive image information relating to a series of at least two images recorded as the vehicle moves along a roadway. The processor is configured to process the image information received by the image receiver to generate an ego-motion estimate of the vehicle, including the translation of the vehicle in the forward direction and the rotation of the vehicle around a vertical axis as between, for example, successive images.

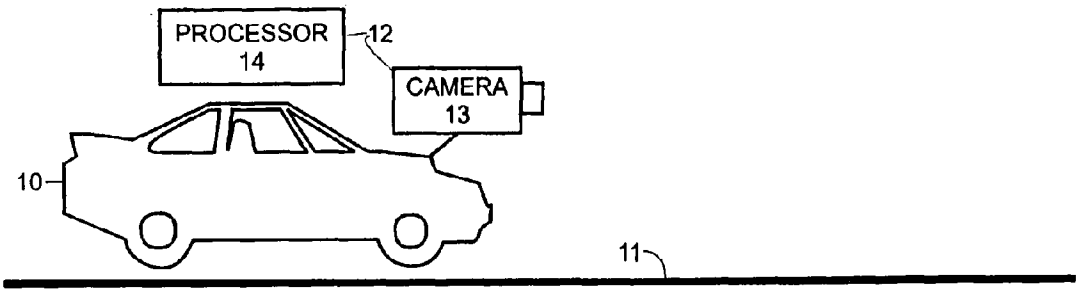


FIG. 1

EP 1 727 089 A3

IPR2025-00943
Tesla EX1002 Page 1552

European Patent Office

EUROPEAN SEARCH REPORT

Application Number
EP 06 01 7695

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
A	US 5 259 040 A (HANNA ET AL) 2 November 1993 (1993-11-02) * abstract; figures 1,2 * * column 4, line 33 - column 6, line 10 *	1-13	INV. G06T7/20 G05D1/02
A	EP 0 390 051 A (HONEYWELL INC) 3 October 1990 (1990-10-03) * abstract; figure 1 * * page 9, line 7 - page 11, line 8 *	1-13	
A	US 5 473 364 A (BURT PETER J [US]) 5 December 1995 (1995-12-05) * column 3, line 61 - column 5, line 19; figures 2,3 *	1-13	

TECHNICAL FIELDS SEARCHED (IPC)
G08G H04N G06T B60K G05D G01S

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
Munich	14 August 2007	Schinnerl, Arnold

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)

IPR2025-00943
Tesla EX1002 Page 1553

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 06 01 7695

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

14-08-2007

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5259040 A	02-11-1993	DE 69231812 D1	07-06-2001
		DE 69231812 T2	15-11-2001
		EP 0606385 A1	20-07-1994
		JP 6511339 T	15-12-1994
		JP 3253962 B2	04-02-2002
		WO 9307585 A1	15-04-1993
EP 0390051 A	03-10-1990	JP 3048977 A	01-03-1991
		US 4969036 A	06-11-1990
US 5473364 A	05-12-1995	JP 10501386 T	03-02-1998
		WO 9534171 A1	14-12-1995

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

IPR2025-00943
Tesla EX1002 Page 1554

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 2 605 185 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
19.06.2013 Bulletin 2013/25

(51) Int Cl.:
G06K 9/00 *(2006.01)*

(21) Application number: **12275201.7**

(22) Date of filing: **12.12.2012**

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR
Designated Extension States:
BA ME

(30) Priority: **12.12.2011 US 201161569306 P**
19.09.2012 US 201261702755 P

(71) Applicant: **Mobileye Technologies Limited**
24854 Nicosia (CY)

(72) Inventor: **Stein, Gideon P**
96263 Jerusalem (IL)

(74) Representative: **Jehan, Robert**
Williams Powell
Staple Court
11 Staple Inn Buildings
London, WC1V 7QH (GB)

(54) Detection of obstacles at night by analysis of shadows

(57) A driver assistance systems mountable in a host vehicle while the host vehicle is moving forward with headlights on for detection of obstacles based on shadows. The driver assistance system includes a camera (12) operatively connectible to a processor (30). A first image frame and a second image frame are captured of a road. A first dark image patch and a second dark image patch include intensity values less than a threshold value. The first and the second dark image patches are tracked from the first image frame to the second image frame as corresponding images of the same portion of the road. Respective thicknesses in vertical image coordinates are measured of the first and second dark image patches responsive to the tracking.

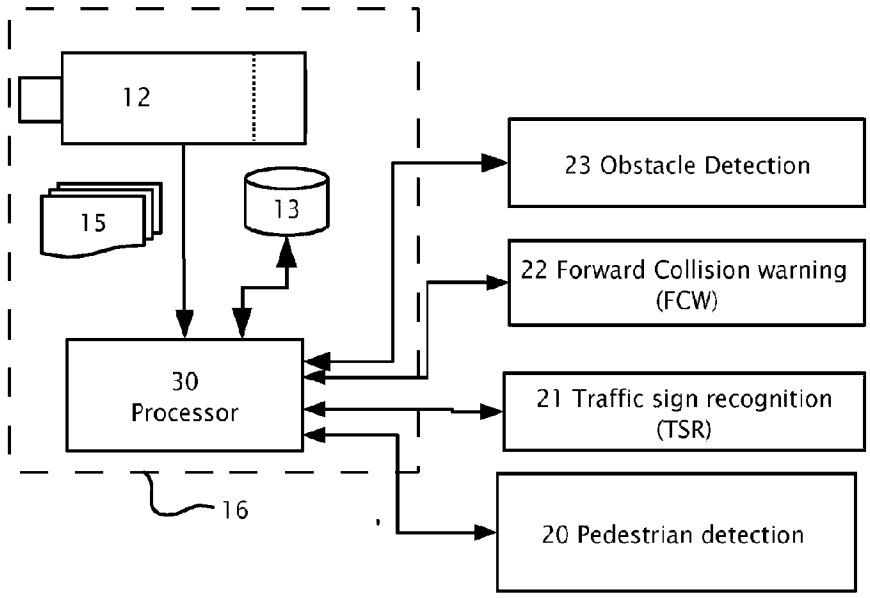


Fig. 2

EP 2 605 185 A2

IPR2025-00943
Tesla EX1002 Page 1555

Description

BACKGROUND

1. Technical Field

[0001] The present invention relates to the field of driver assistance systems (DAS) and in particular to detection of obstacles in a road using a camera.

2. Description of Related Art

[0002] Larson et al. [1] perform detection of negative obstacles using LIDAR and other three dimensional data. Belutta et al. [2] detect depth discontinuities of negative obstacles from stereo range data. Rabkin et al. [3] use stereo range data on thermal cameras because the shape of the terrain also affects the thermal properties. Muraka et al. [4] combine stereo and motion information to detect discontinuities in daytime. They assume a piecewise planar model and show results from indoor and outdoor environment. The range of detection of drop-offs is limited to a few meters since texture is required for stereo and motion correspondence. Witus et al. [5] analyze the shadows produced by light sources mounted on the host vehicle using a stereo camera. Detection range is limited by camera focal length and stereo baseline. Furthermore, depth information on the shadow edge is only available when the light sources are displaced along the stereo baseline. In contrast, a typical stereo camera for driver assistance applications has the two cameras mounted side by side and the headlights are below the cameras.

References

[0003]

1. Larson et al., "Lidar based off-road negative obstacle detection and analysis" In 14th International IEEE Conference on Intelligent Transportation Systems (ITSC), Oct. 2011

2. Belutta et al., "Terrain Perception for Demo III", In Proceedings of the IEEE Intelligent Vehicles Symposium (IV) 2000, October 2000

3. Rabkin et al. "Nighttime negative obstacle detection for off-road autonomous navigation In Unmanned Systems", Technology IX. Edited by Gerhart, Grant R.; Gage, Douglas W.; Shoemaker, Charles M.. Proceedings of the SPIE, Volume 6561, pp. 656103 (2007)

4. Muraka "Detecting obstacles and drop-offs using stereo and motion cues for safe local motion", In IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2008 Sept. 2008

5. Witus et al. Preliminary investigation into the use of stereo illumination to enhance mobile robot terrain perception In Proc. SPIE Vol. 4364, p. 290-301, Unmanned Ground Vehicle Technology III, Grant R. Gerhart; Chuck M. Shoemaker; Eds.]

BRIEF SUMMARY

[0004] The present invention seeks to provide an improved driver assistance system and method.

[0005] According to an aspect of the present invention, there is provided a computerized method performable at night by a driver assistance system mountable in a host vehicle as specified in claim 1.

[0006] According to another aspect of the present invention, there is provided a driver assistance system as specified in claim 11.

[0007] Various driver assistance systems and corresponding methods are provided for herein performable at night for detection of obstacles based on shadows. The driver assistance systems are mountable in a host vehicle while the host vehicle is moving forward with headlights on. The driver assistance systems include a camera operatively connectible to a processor. A first image frame and a second image frame are captured of a road in the field of view of the camera. The first and the second image frames are processed to locate a first dark image patch of the first image frame and a second dark image patch of the second image frame. The first image frame and the second image frame may be filtered with a threshold value of gray scale to produce the first dark image patch of the first image frame and the second dark image patch of the second image frame. The first dark image patch and the second dark image patch include intensity

IPR2025-00943
Tesla EX1002 Page 1556

values less than the threshold value. In another embodiment, the processing to find the first and second dark image patches may include searching for a transition from light to dark (e.g. high gray scale to low gray scale values) and another transition from dark to light. The first and the second dark image patches are tracked from the first image frame to the second image frame as corresponding images of the same portion of the road. Respective thicknesses in vertical image coordinates are measured of the first and second dark image patches responsive to the tracking performed from the first image frame to the second image frame.

[0008] Prior to the tracking, connectivity of picture elements may be performed within the first and second dark image patches. It may be determined that a change of the thicknesses between the first and second dark image patches is consistent with a shadow cast by an obstacle in the road from illumination from the headlights and an obstacle in the road is detected. A driver of the host vehicle may be audibly warned responsive to the detection of the obstacle.

[0009] The change in the thicknesses between the first and second image frames may be analyzed to distinguish between the detected obstacle being a positive or a negative obstacle and/ or a shape of the detected obstacle may be computed.

[0010] It may be determined that a change of the thicknesses between the first and second image frames is not consistent with a shadow cast by an obstacle from illumination from the headlights. The thicknesses may be measured in picture elements according to a total number of picture elements with intensity values less than the threshold divided by a width in picture element columns of the .first and second dark image patches.

[0011] The foregoing and/or other aspects will become apparent from the following detailed description when considered in conjunction with the accompanying drawing figures.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] Embodiments of the present invention are described below, by way of example only, with reference to the accompanying drawings, wherein:

Figures 1 and 2 illustrate a system including a camera or image sensor mounted in a host vehicle, according to an embodiment of the present invention.

Figures 3a and 3b show a simplified line drawings of a night time driving scenario, according to an embodiment of the present invention.

Figures 4a and 4b show image frames of a speed bump where a host vehicle is at different distances from the speed bump to illustrate an embodiment of the present invention.

Figures 4c and 4d show a portion of image frames as shown respectively in Figures 4a and 4b in greater detail to illustrate an embodiment of the present invention.

Figure 5a shows an example of a shadow produced , according to an embodiment of the present invention.

Figures 5b, 5c and 5d show further details of the drop in road surface as shown in Figure 5a, at different distances to the drop.

Figure 6a shows a change in asphalt color with no drop, according to an embodiment of the present invention.

Figure 6b shows a gray-scale profile for a simple step in brightness with no dark shadow patch for the change in asphalt color shown in Figure 6a.

Figures 6c and 6e show a dark stripe from a change in asphalt paving of the road shown in successive image frames, according to an embodiment of the present invention.

Figures 6d and 6f show a simple step edge of brightness in greater detail of respective

Figures 6c and 6e, according to an embodiment of the present invention.

Figure 7 which shows a graph of simulation results, according to an embodiment of the present invention.

Figures 8a and 8b shows a simplified line drawing of a night time driving scenario and an area of greater detail of the night time driving scenario respectively, according to an embodiment of the present invention.

Figure 9 shows a graph of a simulation for a drop in road surface of 0.1m over a length of 0.8m and 1.6m, according to an embodiment of the present invention.

Figures 10a-10e show a sequence of five images where a host vehicle approaches a speed bump and a shadow, according to an embodiment of the present invention.

Figures 10f-10j show greater detail of the speed bump and the shadow for respective Figures 10a-10e.

Figure 11 shows a flow chart of a method, according to an embodiment of the present invention.

Figure 12 shows a graph of the thickness of shadow at the far side of a speed bump as a function of distance to the speed bump, according to an embodiment of the present invention.

Figures 13a and 13b show two examples of soft shoulders at the road edge, according to an embodiment of the present invention.

IPR2025-00943
Tesla EX1002 Page 1557

DETAILED DESCRIPTION

[0013] Reference will now be made in detail to embodiments of the present invention, examples of which are illustrated in the accompanying drawings, wherein like reference numerals refer to the like elements throughout. The features are described below to explain the teachings herein by referring to the Figures.

[0014] Before explaining features taught herein in detail, it is to be understood that the invention is not limited in its application to the details of design and the arrangement of the components set forth in the following description or illustrated in the drawings. The invention is capable of other features or of being practiced or carried out in various ways. Also, it is to be understood that the phraseology and terminology employed herein is for the purpose of description and should not be regarded as limiting.

[0015] By way of introduction, embodiments of the present invention are directed to driver assistance systems using forward facing cameras which are becoming mainstream. These forward facing cameras may be rigidly mounted inside the driver cabin near the rear-view mirror. Cameras mounted inside the driver cabin near the rear-view mirror puts the cameras at a fixed position, significantly above the vehicle headlights. With the fixed position, the shadows produced by the headlights, when driving on dark roads, are visible in the camera image. The following detailed description describes how to use these shadows to enhance obstacle detection. In particular the description describes how the shadows can be used for detecting speed-bumps and negative obstacles at night from a moving vehicle. The information from the enhanced obstacle detection can be used to warn the driver or to prepare the vehicle for the anticipated change in road surface. The detailed description also gives a detailed analysis of the particular image motion characteristics of the shadow when the host vehicle is moving forward. The information from the detailed analysis allows the discrimination between shadows and other dark patches on the road and then, through careful analysis of the shadow over time, both the depth and shape of the obstacle can be determined. The detection range from shadow analysis may be significantly larger than what is possible from stereo or structure from motion (SFM) in dark night conditions.

[0016] As host vehicle moves forward, the image of a dark patch of a shadow, produced by the host vehicle headlights, behaves very differently to other dark patches on the road such as those due to tar patches, differences in asphalt color or shadows due to other light sources.

[0017] It will be shown below that the image of a mark on a road surface will increase as a function of the inverse of the distance (Z) squared, where Z is the distance from the camera to the mark on the road surface:

$$\frac{1}{Z^2} = Z^{-2}$$

[0018] While the size of a shadow produced by a sharp negative edge will increase only as a function of the inverse of the distance (Z):

$$\frac{1}{Z} = Z^{-1}$$

[0019] The shadow of a round edge, such as the far side of a speed bump, will increase even more slowly. By tracking the dark patch over time it is possible to differentiate between shadows and road marks.

[0020] The difference in the observed behavior can be explained as follows. A mark on the road retains its size (in world coordinates) as a host vehicle approaches. With a sharp negative obstacle, the visible part of the shadow extends from the top of the negative obstacle step to the far side of the shadow. When the host vehicle moves towards the obstacle, the lighting angle changes and the far edge of the shadow on the road moves and the extent of the shadow on the road is decreased significantly. If the edge producing the shadow is rounded the closer edge of the shadow is formed by the tangent point of the edge. As the host vehicle moves forward the tangent point moves further away and drops in height. These two factors of the edge and the host vehicle moving forward, reduce the size of the shadow on the road even further.

[0021] Referring now to the drawings, reference is now made to Figures 1 and 2 which illustrate a system **16** including a camera or image sensor **12** mounted in a host vehicle **18**, according to an aspect of the present invention. Image sensor **12**, imaging a field of view in the forward direction provides image frames **15** in real time and image frames **15** are captured by an image processor **30**. Processor **30** may be used to process image frames **15** simultaneously and/or in parallel to serve a number of advanced driver assistance systems/applications. The advanced driver assistance systems (ADAS) may be implemented using specific hardware circuitry with on board software and/or software control

IPR2025-00943
Tesla EX1002 Page 1558

algorithms in memory **302.** Image sensor **12** may be monochrome or black-white, i.e. without color separation or image sensor **12** may be color sensitive. By way of example in Figure 2, image frames **15** are used to serve pedestrian detection **20,** traffic sign recognition (TSR) **21** forward collision warning (FCW) **22** and obstacle detection **23** using shadows, according to embodiments of the present invention. Processor **30** may be used to process image frames **15** to detect and recognize an image or portions of the image in the forward field of view of camera **12.**

[0022] In some cases, image frames **15** are partitioned between different driver assistance applications and in other cases the image frames **15** may be shared between the different driver assistance applications.

[0023] Although embodiments of the present invention are presented in the context of driver assistance applications, embodiments of the present invention may be equally applicable in other real time signal processing applications and/or digital processing applications, such as communications, machine vision, audio and/or speech processing as examples.

[0024] Reference is now made to Figure 3a which shows a simplified line drawing **30a** of a night time driving scenario, according to a feature of the present invention. Camera **12** is mounted in close proximity to the host vehicle **18** windshield and host vehicle **18** headlights illuminate the scene ahead as host vehicle **18** travels along road surface **36.** Host vehicle **18** headlights produce a shadow **32** on the road when there is a obstacle. In the example shown in Figure 3a the obstacle is low object **34** on road **36.** Low object **34** may be for example a speed bump, raised manhole cover but the discussion that follows is the same for a negative obstacle such as a drop or pothole. Since the camera **12** is mounted higher than the headlights, shadow **32** can be seen by camera **12** in host vehicle **18.** Negative obstacles and speed bumps produce dark patches in image frames **15.**

[0025] In Figure 3a are a number of distances, where distance is designated as Z in the equations which are shown later on the detailed descriptions that follow. Let the coordinate system reference point (0,0) be the point on the road **36** directly below camera **12,** indicated by dotted line **31.** H_C is the height of camera **12,** H_L the height of the headlight of host vehicle **18** and H_T the height of obstacle **34.** The headlights of vehicle **18** are Z_L in front of camera **12,** that is the distance between dotted lines **31** and **33.** Z_T is the distance from camera **12** to the edge of obstacle **34,** that is the distance between dotted lines **31** and **35.** Z_{S1} and Z_{S2} are the close and far points on shadow **32** which are visible to camera **12,** the distances between dotted lines **31** and **37** and dotted lines **31** and **39** respectively.

[0026] Reference is now made to Figure 3b which shows a simplified line drawing **30b** of a night time driving scenario, according to an embodiment of the present invention. Two host vehicles **18** are shown with respective cameras **12.** From Figure 3b it can be seen that when host vehicle **18** moves forward towards obstacle **34,** the extent of shadow **32** becomes smaller.

[0027] When host vehicle **18** moves towards the obstacle **34,** the lighting angle changes and the far edge of shadow **32** on road **36** moves and the extent of shadow **32** on road **36** is decreased significantly. If the edge producing shadow **32** is rounded, the closer edge of shadow **32** is formed by the tangent point of the edge. As host vehicle **18** moves forward the tangent point moves further away and drops in height. These two factors of the edge and the host vehicle **18** moving forward, reduce the size of shadow **32** on road **36** even further.

[0028] Reference is now made to Figures 4a and 4b which show image frames of a speed bump where host vehicle **18** is at a distance away from the speed bump of 40.9m and 27.0m respectively, according to an embodiment of the present invention. Figures 4a and 4b show white reflectors **40** and a rock **42.** The shadow beyond the speed bump is only slightly wider in Figure 4b than the shadow in Figure 4a where host vehicle **18** is farther from the speed bump. Also, the lateral distance difference between reflectors **40** is much more pronounced when compared to the change in shadow width beyond the speed bump for the two image frames. Therefore, the shadow on the far side of the speed bump increases and/ or decreases in size much more slowly than the image size of the speed bump itself.

[0029] Reference is now made to Figures 4c and 4d which show a portion of image frames shown respectively in Figures 4a and 4b in greater detail, according to an embodiment of the present invention. Plotted on each of Figures 4a and 4b are respective gray scale value curves **44a** and **44b.** The further to right curves **44a/ 44b** go laterally, the greater is the gray scale value. Conversely, the further to left curves **44a/ 44b** go laterally, the lesser is the gray scale value. In Figure 4c where host vehicle **18** is farther from the speed bump, peaks in gray scale value occur by virtue of the contributions from reflection of rock **42** and the illuminated side of the speed bump along with reflection from reflectors **40.** Furthermore, it can be observed that the shadow cast by the speed bump is only slightly narrower when the host vehicle **18** is farther from the speed bump.

[0030] Changes in road **36** surface material, ice patches, oil marks and tar seams can also produce dark patches on the image of the road **36** surface. However, the size of such a patch in the image **15,** as the host vehicle **18** approaches, behaves very differently.

[0031] Reference is now made to Figure 5a shows an example of a shadow produced when there is a drop **52** in the road surface where there is a junction between older and newer asphalt paving. Reference is also made to Figures 5b, 5c and 5d shows further details of drop **52,** when the distance to drop **52** is at 22meters (m), 17.7m and 11.5m respectively. In Figures 5b-5d it can be seen that the size of the shadow of drop **52,** where newer paving makes way to old paving and when the road is illuminated by the headlights of host vehicle **18.** The shadow of drop **52** hardly changes as host vehicle **18** distance (Z) to drop **52** changes from 22m to 11.5m. Figures 5b, 5c and 5d have respective gray scale value

IPR2025-00943
Tesla EX1002 Page 1559

curves **54b**, **54c** and **54d** for a central column of pixels. The lowest gray scale value is shown by the left lateral peak in respective curves **54b**, **54c** and **54d** in the center of the shadow of drop **52.**

[0032] Reference is now made to Figure 6a which shows a change in asphalt color with no drop , according to an embodiment of the present invention. The corresponding gray-scale profile for Figure 6a is shown in Figure 6b, which shows a simple step in brightness with no dark shadow patch.

[0033] Reference is now made to Figures 6c and 6e which show a dark stripe **62** which is a change in asphalt paving of the road shown in successive image frames **15,** according to an embodiment of the present invention. The change in asphalt paving of the road is due to a double change in asphalt by virtue of a ditch which was filled in with new asphalt. In each image frame **15** their is no drop in height of the road surface and no shadow is produced. The gray-scale profiles shown in Figures 6d and 6f show a simple step edge of brightness in greater detail of respective Figures 6c and 6e, according to an embodiment of the present invention. Figures 6c and 6e show a dark stripe **62** on the road at 14.3m and 11.2m. The width of dark stripe **62** as shown in Figures 6c and 6e increase in size as a function of the inverse of distance squared (Z^{-2}) indicating that dark stripe **62** is a mark on the road and not a shadow from a negative obstacle (obstacle **34** for example).

[0034] The techniques described below, use the change in the actual shadow **32** itself due to the host vehicle **18** motion (and thus the motion of the light sources). The actual shadow is then analyzed by the camera **12** that is fixed relative to the light source and moves with it.

THE MATHEMATICAL MODEL

[0035] Referring back to Figure 3a, reference point (0,0) is the point on the road **32** directly below camera **12.** H_C is the height of the camera, H_L the height of the headlight and H_T the height of obstacle **34.** The headlights are Z_L in front of the camera. Z_T is the distance from camera **12** to the edge of obstacle **34** (dotted line **35)** and Z_{S1} and Z_{S2} are the close and far points of shadow **32** visible to camera **12.**

[0036] By similarity of triangles we get:

$$Z_{S1} = \frac{Z_T H_C}{H_C - H_T} \quad (1)$$

$$Z_{S2} = \frac{(Z_T - Z_L)H_L}{H_L - H_T} + Z_L \quad (2)$$

$$= \frac{Z_T H_L - Z_L H_T}{H_L - H_T} \quad (3)$$

[0037] Let y1 and y2 be the lower and upper edges of a shadow patch in the image:

$$y_1 = \frac{f H_C}{Z_{S1}} \quad (4)$$

$$y_2 = \frac{f H_C}{Z_{S2}} \quad (5)$$

[0038] The height or thickness of the patch in the image is then given by y1 -y2:

IPR2025-00943
Tesla EX1002 Page 1560

$$y_1 - y_2 = \frac{fH_C}{Z_{S1}} - \frac{fH_C}{Z_{S2}} \quad (6)$$

$$= \frac{fH_C(H_C - H_T)}{Z_T H_C} - \frac{fH_C(H_L - H_T)}{Z_T H_L - Z_L H_T} \quad (7)$$

[0039] For $Z_L << Z_T$ we can approximate:

$$y_1 - y_2 \approx \frac{f(H_T H_C - H_T H_L)Z_T}{Z_T^2 H_C} \quad (8)$$

$$\approx \frac{fH_T(H_C - H_L)}{Z_T H_C} \quad (9)$$

[0040] The image thickness of the shadow patch changes as a function of Z_T^{-1} and given an estimate of Z_T one can also estimate the (negative) height of obstacle **34.**

[0041] Consider now a marking on road **36** which at time to stretches from a distance Z_{S1} (to) to a distance Z_{S2} (to). Let:

$$\alpha = Z_{S2}(t_0) - Z_{S1}(t_0) \quad (10)$$

$$Z_{S3} = Z_{S1} + \alpha \quad (11)$$

[0042] The height or thickness of the marking in the image is given by:

$$y_3 - y_1 = \frac{fH_C}{Z_{S1} + \alpha} - \frac{fH_C}{Z_{S1}} \quad (12)$$

$$= \frac{fH_C\alpha}{(Z_{S1} + \alpha)Z_{S1}} \quad (13)$$

$$\approx \frac{fH_C\alpha}{Z_{S1}^2} \quad (14)$$

NUMERICAL EXAMPLE

[0043] Reference is now made to Figure 7 which shows a graph **70** for a simulation, according to an embodiment of the present invention. Graph **70** plots two curves, one for a patch or road mark (dotted line) and one for a shadow strip (solid line). The thickness of the patch or shadow strip is in pixels versus range to obstacle in meters (m). The simulation

IPR2025-00943
Tesla EX1002 Page 1561

has camera **12** mounted in a Renault Megane II where H_C = 1.25m, H_L = 0.73m and obstacle of height H_T = 0.1m. A shadow will appear just over 2 pixels thick at 50m. At 25m the shadow will double in size to about 4 pixels. A road mark on road at 25m will quadruple in size to 8 pixels.

[0044] Note that at a certain close distance (≈ 5m) the tangent line of the headlight and of the camera **12** coincide. After passing 5m the size of the shadow patch decreases rapidly.

[0045] Referring back to Figures 3a and 3b, Figures 3a and 3b show a step edge drop. The step edge drop is typical for a pot hole, road edge and steep speed bump. With the step edge drop, the tangent point which causes shadow **32** and the tangent point of the viewing angle from camera **12** can be approximated as the same point and which is fixed in space as host vehicle **18** moves. For extended smooth speed bumps this approximation breaks down, however.

[0046] Reference is now made to Figures 8a and 8b which shows a simplified line drawing **80** of a night time driving scenario and an area **83** of greater detail of the night time driving scenario respectively, according to an embodiment of the present invention. In Figures 8a and 8b, the tangent viewing ray **85** touches the speed bump **34** further away than the tangent illumination ray **87.** As host vehicle **18** moves forward, the tangent point **62** of the viewing ray **65** moves further away along the speed bump or obstacle **34.** This means that the near edge of shadow **32,** in image space, will move slower than a typical road point. The exact motion depends on the shape of speed bump **34** and can be used to measure the curvature of obstacle **34.**

[0047] It can also be seen that as host vehicle **18** moves forward, tangent point **84** of illumination ray **87** moves further away along obstacle **34** and the height H_T decreases. The result is that the far edge of shadow **32** on road **36** will move towards the host vehicle **18** even faster than predicted by a sharp edge. Shadow **32** on road **36** moving towards the host vehicle **18** even faster than predicted by the sharp edge effect will accelerate when host vehicle **18** gets close to the speed bump or obstacle **34** and at some distance shadow **32** will disappear completely.

[0048] From experimental results shown later, at far distances (above 30m), shadow **32** behaves more like a step edge and disappears at about 7.0m from camera **12.**

[0049] The analysis of rounded edges is complex. The tangent point **84** for the headlights does not correspond exactly to the tangent point **82** of camera **12** line of sight. Furthermore, as the host vehicle **18** moves closer to obstacle **34,** tangent points **82** and **84** will move further away and closer to road **36.** In order to analyze a rounded edge obstacle **34** (e.g. speed bump), the edge was approximated as a piecewise linear curve of 10 segments dropping from H_T = 0.1m down to zero over a given length:

$$Z_T(i) = Z_T(0) + i * \delta Z \qquad (15)$$

$$H_T(i) = H_T * (1 - 0.1 * i) \qquad (16)$$

where i=1...10.

[0050] Reference is now made to Figure 9 which shows a graph **90** for a simulation for a drop of 0.1 meter in road surface over a length of 0.8 meter and 1.6 meter, according to an example. Three curves are shown for thickness (pixels) of shadow versus range to speed bump or obstacle **34** in meters (m). Solid line **92** is for a sharp edge, dotted line **94** is for a round edge for the drop of 0.1m over the length of 0.8m and dotted line **96** is for another round edge for the drop of 0.1m over the length of 1.6m.

[0051] Host vehicle **18** in the simulation, approaches the speed bump or obstacle **34** from 50 meter down to 5 meter in steps of 1 meter. At each distance it was determined (for each of the two cases: 0.8m and 1.6m) which point along the piecewise linear curve was the occluding point for the headlight and produced the far edge and which was the occluding point for camera **12** and produced the near edge. As can be expected, the shorter the drop, the more it behaves like a step edge.

EXPERIMENTS AND RESULTS

[0052] Figures 10a-10e show a sequence of five image frames **15** where host vehicle **18** approaches speed bump **1002** and shadow **1006,** according to an embodiment of the present invention. Figures 10f-10j show greater detail of speed bump **1002** and shadow **1006** for respective Figures 10a-10e. In Figures 10a-10e the distance (Z) to speed bump **1002** and shadow **1006** are 66m, 42.4m, 26.4m, 13.6m and 7.9m respectively. Shadow **1006** at the far side of speed bump **1002** was first detected at over 40m and then tracked as it grew closer to host vehicle **18.** Figures 10f-10j have respective gray scale value curves **1004f-1004j** for a central column of pixels. The lowest gray scale value for the central column is shown by the lateral left peak in respective curves **1004f-1004j** in the center of shadow **1006.**

IPR2025-00943
Tesla EX1002 Page 1562

[0053] Reference is now made to Figure 11 which shows a flow chart of a method **23,** according to an embodiment of the present invention.

[0054] In step **1103,** a first image frame **15** is captured, followed by the capture of a second image frame **15** from camera **12** processor **30** of system **16.**

[0055] Optionally, image frames **15** may be warped to compensate for a change in vehicle orientation, i.e. roll, pitch and yaw, which may be determined from image frames **15** or from external sensor(s) connected to driver assistance system **16.**

[0056] The first and the second image frames **15** are then filtered (step **1105)** according to a threshold, e.g. a previously determined threshold, to produce a first dark image patch **1107** of first image frame **15** and a second dark image patch **1111** of the second image frame **15** respectively. First image patch **1107** and optionally second image patch **1111** are analyzed for connectivity of picture elements.

[0057] One or more connected components of first image patch **1107** is tracked (step **1112)** from the first image frame **15** to second dark image patch **1111** of second image frame **15.**

[0058] Tracking step **1113** may be performed by using the expected image motion for road points given host vehicle **18** speed and nearest neighbor.

[0059] The thicknesses of first image patch **1107** and second dark image patch **1111** are measured (step **1113)** in both the first image frame **15** and the second image frame **15.** By way of example, an image strip of twenty pixel columns wide, from the bottom of an image frame **15** and up to the image row corresponding to 80 meters may be analyzed for connected components. The patch or shadow thickness may be computed (step **1113)** by counting the number of pixels in the connected component divided by the strip width in columns (i.e. 20 pixels).

[0060] In decision block **1115,** the thicknesses of candidate shadows **1109** tracked across successive image frames may be accumulated. From the change of thicknesses tracked across successive image frames for first dark image patch **1107** and second dark image patch **1111** it is possible to analyze (step **1117)** if the change of thicknesses is consistent with a shadow cast by an obstacle from illumination from the headlights in order to for instance audibly warn a driver of host vehicle **18** of an obstacle **34** in road **36,** or to apply brakes without driver intervention. First image patch **1107** and second dark image patch **1111** may be shadows of an obstacle from the headlights of vehicle **18** or alternatively a darker road texture unrelated to an obstacle in the road.

[0061] Reference is now made to Figure 12 which shows a graph **1200** of the thickness of shadow **1006** at the far side of speed bump **1002** as a function of distance to speed bump **1002,** according to feature of the present invention. Dots show the results measured in the sequence of frames captured. The curves **1202** (sharp edge) and **1204** (round edge) shown the simulation of a bump 0.1m high for a sharp drop and a smooth drop respectively. The results show that in practice the shadows behave according to the model. Shadow 1006 behaves approximately like Z^{-1} and not Z^{-2}. Shadow **1006** appears to behave like a rounded drop however, further experiments may be required using ground truth from accurately measured road structures.

[0062] Reference is now made to Figures 13a and 13b which show two examples of soft shoulders at the road edge, according to an example of the present invention. The drop between the asphalt and the gravel shoulder, results in a narrow shadow line **1300a** and **1300b** respectively that are darker than both the asphalt or the gravel. Often in country roads the asphalt paved on top of gravel and there is a step transition (drop) from the asphalt to the gravel. The step transition is called a soft shoulder which is a negative obstacle and will also generate a shadow like shadow lines **1300a** and **1300b.** Soft shoulders can thus be detected by looking for a dark strip a few pixels wide at the edge between the dark asphalt and the lighter gravel. The shadow strip will be even darker than the asphalt. The width of shadow lines **1300a** and **1300b** indicate the shoulder drop. These dark shadow lines **1300a** and **1300b** have a unique feature; unlike road markings such as lane marks that get narrower in the image as they approach the horizon, the shadows like **1300a** and **1300b** on road edges stay the same width. On a straight road, the inner edge of the shadow would be a line passing through the vanishing point of the road while the outer edge is a line that passes a few pixels to the side of this vanishing point.

[0063] The term "obstacle" as used herein in reference to a road refers to a "positive obstacle" and a "negative obstacle". The term "positive obstacle" as used herein is an obstacle extends vertically upward in real space above the road surface such as a speed bump or a vertical bump caused by new asphalt paving in the direction of motion of the vehicle. The term "negative obstacle" as used herein extends vertically downward in real space below the road surface such as a hole or a drop in asphalt paving in the direction of motion of the vehicle.

[0064] The term "shape" of an obstacle, as used herein refers to the vertical contour in real space along the direction of motion of the vehicle

[0065] The indefinite articles "a", "an" is used herein, such as "a shadow", an "obstacle" has the meaning of "one or more" that is "one or more shadows" or "one or more obstacles". Although selected embodiments of the present invention have been shown and described, it is to be understood the present invention is not limited to the described features. Instead, it is to be appreciated that changes may be made to these features without departing from the principles and spirit of the invention, the scope of which is defined by the claims and the equivalents thereof.

IPR2025-00943
Tesla EX1002 Page 1563

[0066] The features of the various embodiments described herein may be combined with one another.

[0067] The disclosures in United States patent application numbers US-61/569,306 and US-61/702,755, from which this application claims priority, and in the abstract accompanying this application are incorporated herein by reference.

Claims

1. A computerized method performable at night by a driver assistance system mountable in a host vehicle while the host vehicle is moving forward with headlights on, wherein the driver assistance system includes a camera operatively connectible to a processor, the method including the steps of:

   capturing a first image frame and a second image frame of a road in the field of view of the camera;
   processing said first image frame and said second image frame to locate a first dark image patch of said first image frame and a second dark image patch of said second image frame;
   tracking said first and said second dark image patches from said first image frame to said second image frame as corresponding images of the same portion of the road; and
   measuring respective thicknesses in vertical image coordinates of said first and second dark image patches responsive to said tracking from said first image frame to said second image frame.

2. The method of claim 1, wherein said processing is performed by filtering said first image frame and said second image frame with a threshold value of gray scale to produce said first dark image patch of said first image frame and said second dark image patch of said second image frame, wherein said first dark image patch and said second dark image patch include intensity values less than said threshold value;

3. The method of claim 1 or 2, including the step of:

   prior to said tracking, verifying connectivity of picture elements within said first dark image patch.

4. The method of claim 1, 2 or 3, including the step of:

   determining that a change of said thicknesses between said first and second dark image patches is consistent with a shadow cast by an obstacle in the road from illumination from the headlights, thereby detecting said obstacle in the road.

5. The method of claim 4, including the step of:

   audibly warning a driver of the host vehicle responsive to said detecting said obstacle.

6. The method of claim 4 or 5, including the step of:

   analyzing change in said thicknesses between said first and second image frames, thereby distinguishing between the detected obstacle being a positive or a negative obstacle.

7. The method of claim 4, 5 or 6, including the step of:

   analyzing change in said thicknesses between said first and second image frames, thereby computing a shape of the detected obstacle.

8. The method of any preceding claim, including the step of:

   determining that a change of said thicknesses between said first and second image frames is not consistent with a shadow cast by an obstacle from illumination from the headlights.

9. The method of any preceding claim, including wherein said thicknesses are measured in picture elements according to a total number of picture elements with intensity values less than said threshold divided by a width in picture element columns of said .first and second dark image patches.

10. A driver assistance system including a camera and a processor configured to perform the computerized method,

IPR2025-00943
Tesla EX1002 Page 1564

according to any preceding claim.

11. A driver assistance system mountable in a host vehicle, the driver assistance system includes a camera operatively connectible to a processor, wherein the driver assistance system while the host vehicle is moving is operable to:

capture a first image frame and a second image frame of a road in the field of view of the camera;
process said first image frame and said second image frame to locate a first dark image patch of said first image frame and a second dark image patch of said second image frame;
track said first and second dark image patches from said first image frame to said second image portion of said second image frame as corresponding images of the same portion of the road; and
measure respective thicknesses in vertical image coordinates of said first and second dark image patches responsive to said tracking from said first image frame to said second image frame.

12. The driver assistance system of claim 11, wherein the system is operable to determine that a change of said thicknesses between said first and second dark image patches is consistent with a shadow cast by an obstacle in the road from illumination from the headlights, thereby detecting said obstacle in the road.

13. The driver assistance system of claim 12, wherein the system is operable to analyze change in said thicknesses between said first and second image frames to distinguish between the detected obstacle being a positive or a negative obstacle.

14. The driver assistance system of claim 12 or 13, wherein the system is operable to analyze change in said thicknesses between said first and second image frames to compute a shape of the detected obstacle.

15. The driver assistance system of any one of claims 11 to 14, wherein the system is operable to determine that a change of said thicknesses between said first and second image frames is not consistent with a shadow cast by an obstacle from illumination from the headlights.

IPR2025-00943
Tesla EX1002 Page 1565

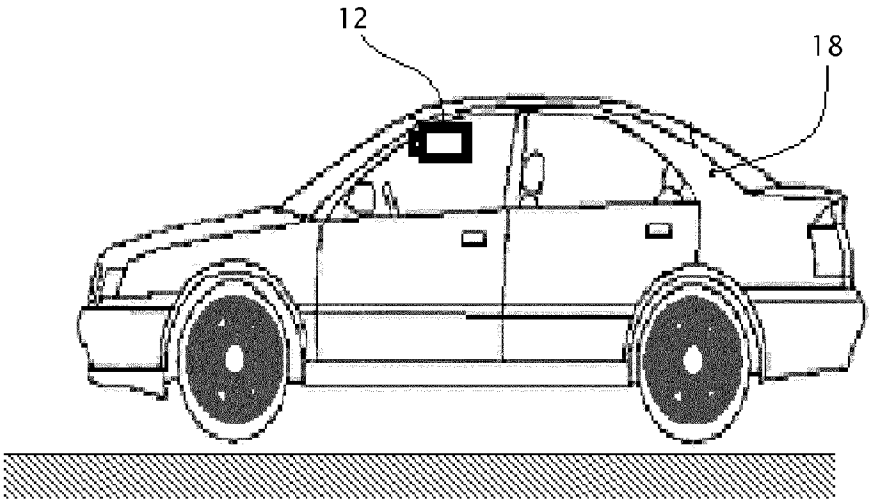


Fig. 1

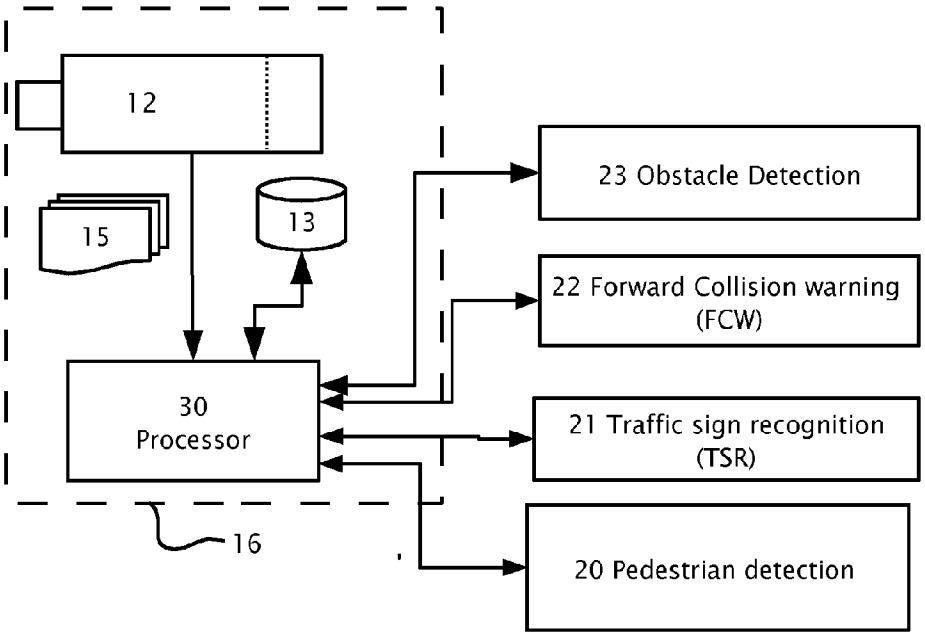


Fig. 2

IPR2025-00943
Tesla EX1002 Page 1566

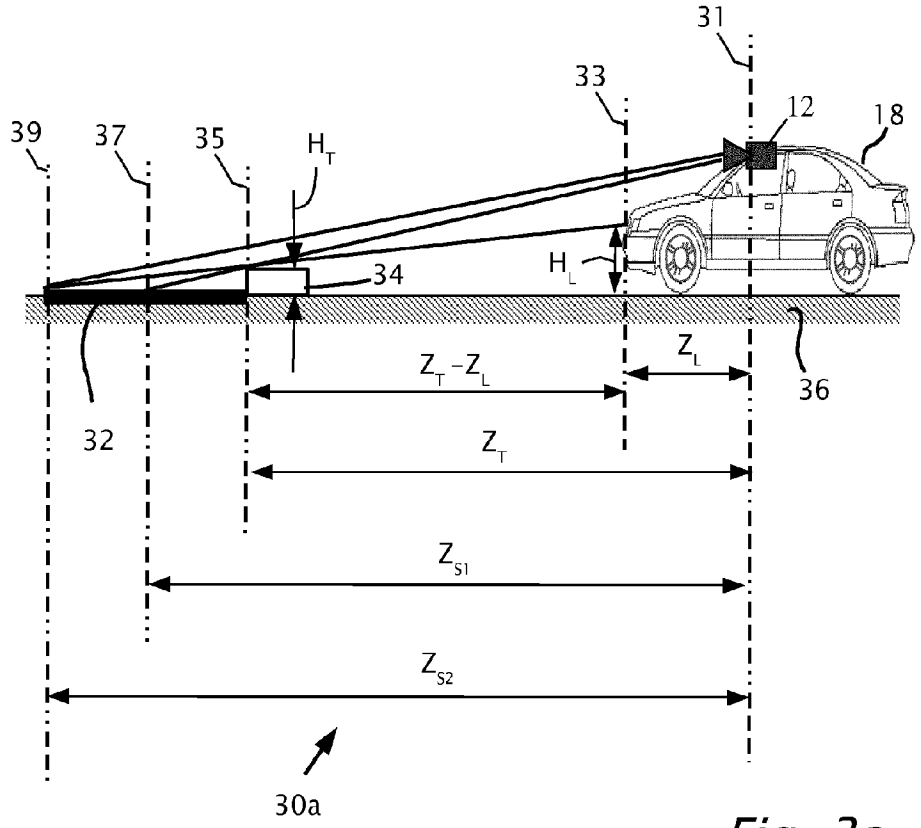


Fig. 3a

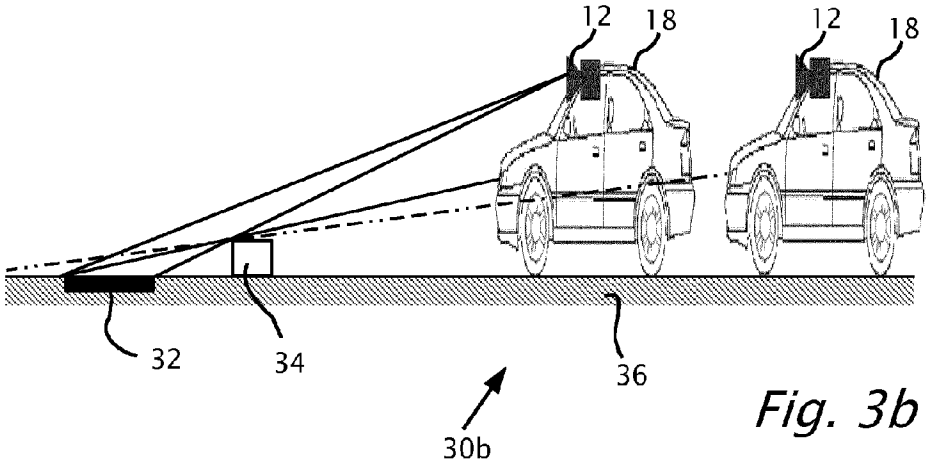


Fig. 3b

IPR2025-00943
Tesla EX1002 Page 1567

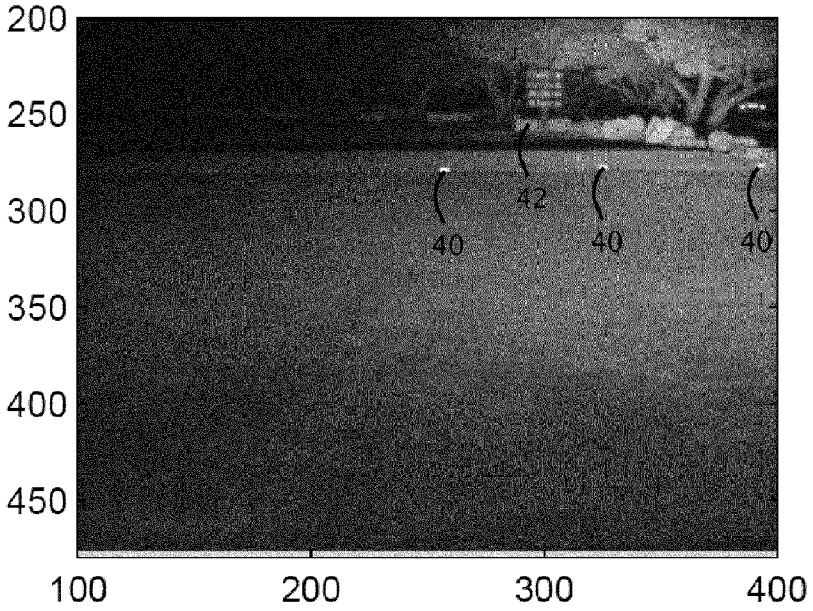


Fig. 4a

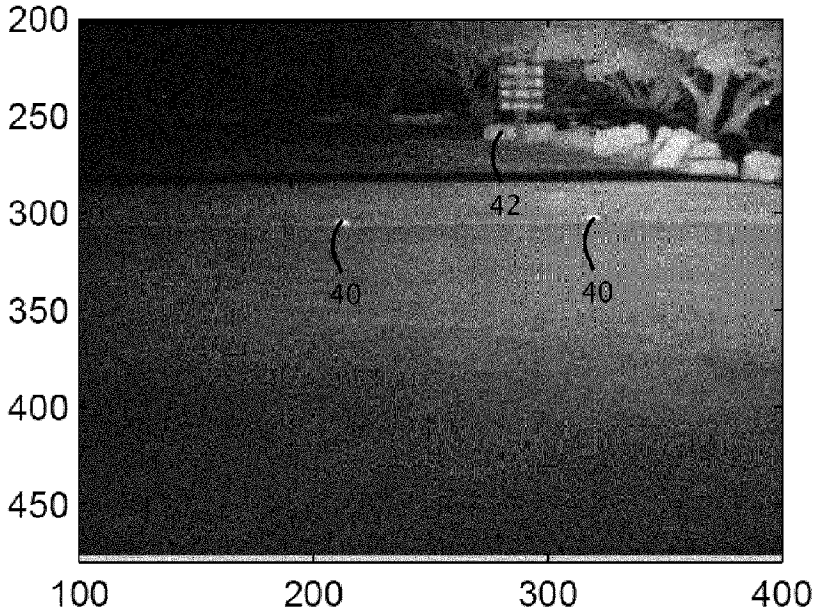


Fig. 4b

IPR2025-00943
Tesla EX1002 Page 1568

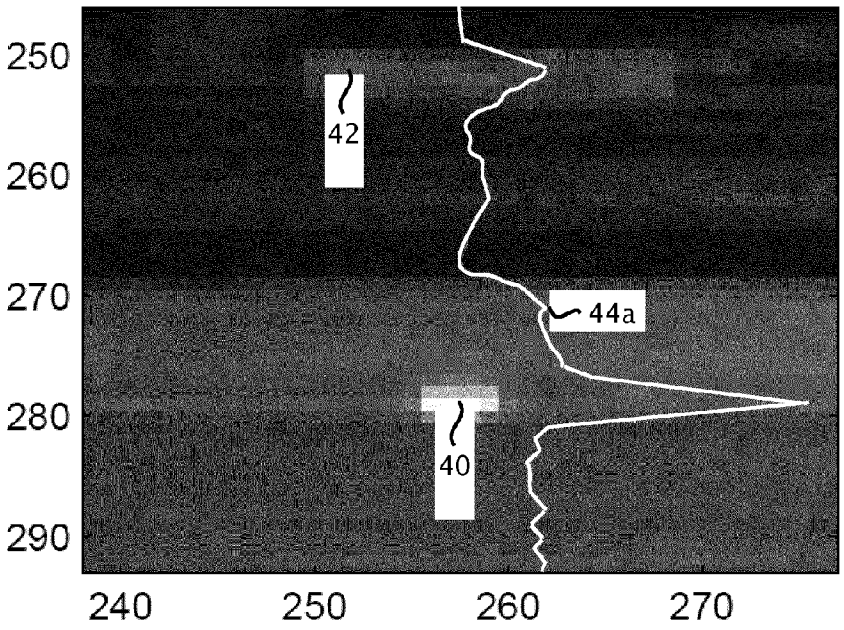


Fig. 4c

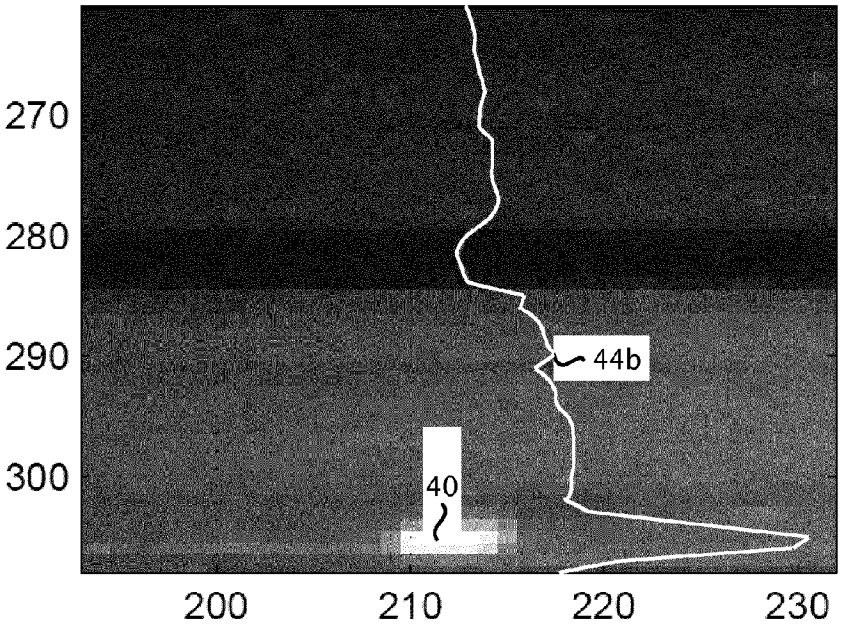


Fig. 4d

IPR2025-00943
Tesla EX1002 Page 1569

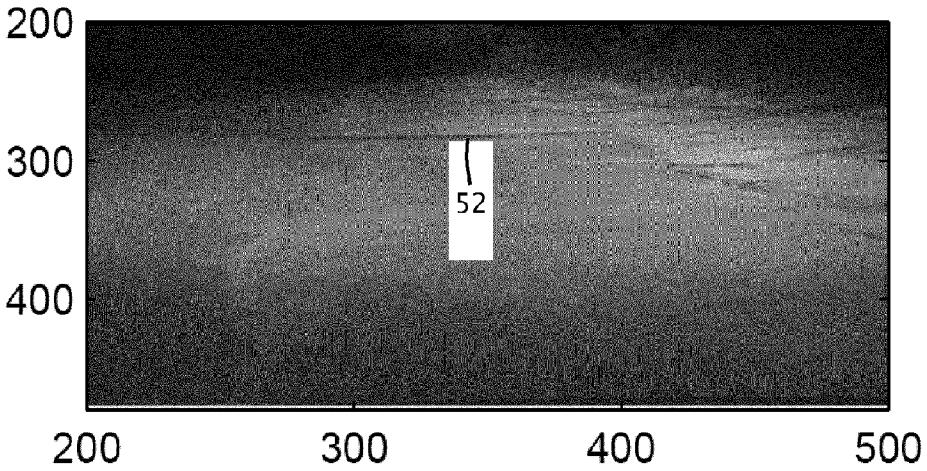


Fig. 5a

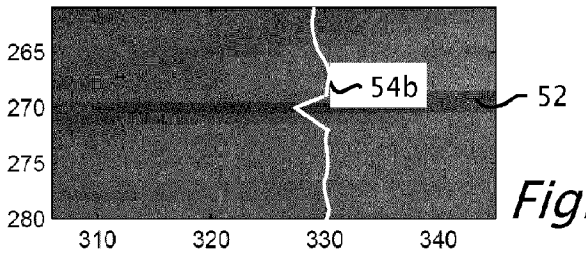


Fig. 5b

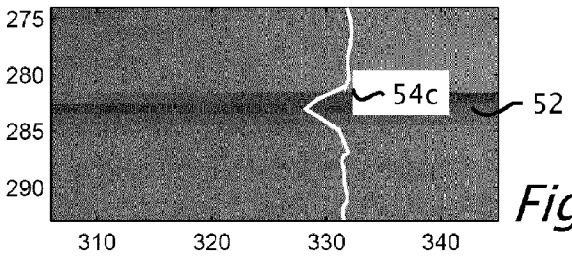


Fig. 5c

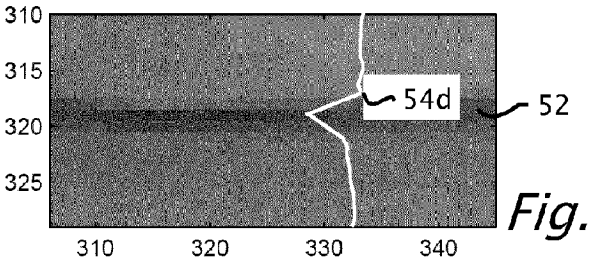


Fig. 5d

IPR2025-00943
Tesla EX1002 Page 1570

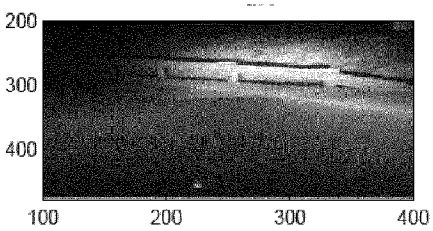


Fig. 6a

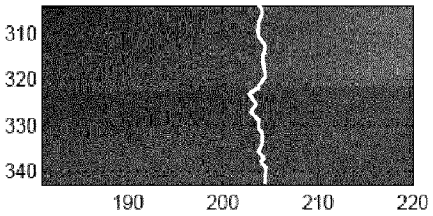


Fig. 6b

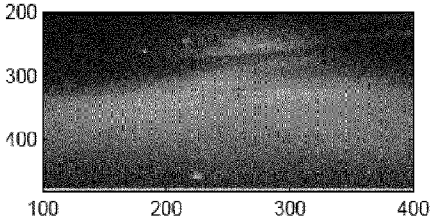


Fig. 6c

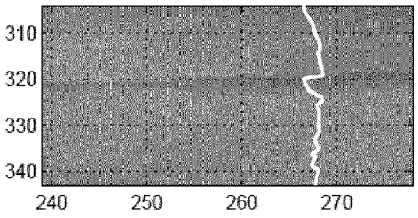


Fig. 6d

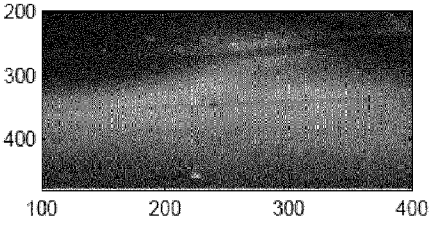


Fig. 6e

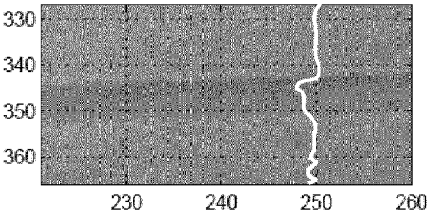


Fig. 6f

IPR2025-00943
Tesla EX1002 Page 1571

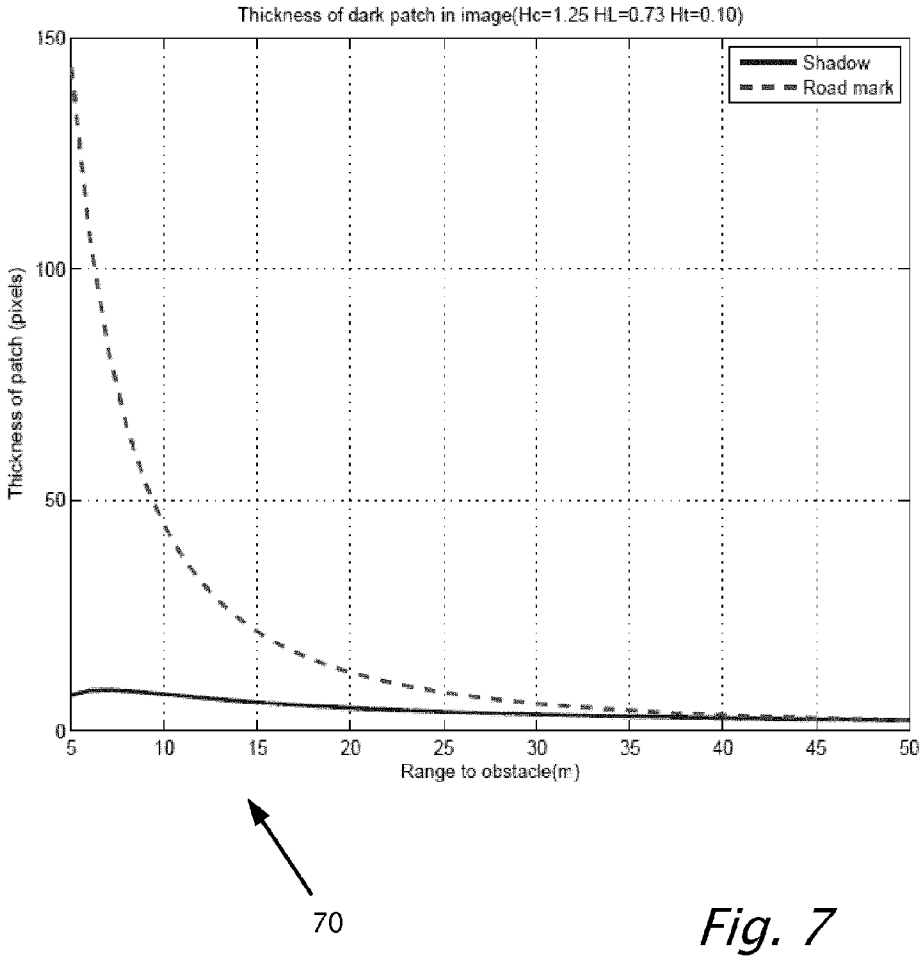


Fig. 7

IPR2025-00943
Tesla EX1002 Page 1572

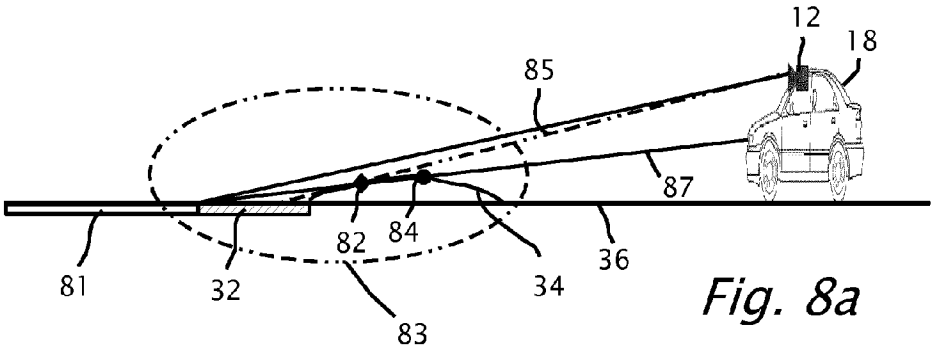


Fig. 8a

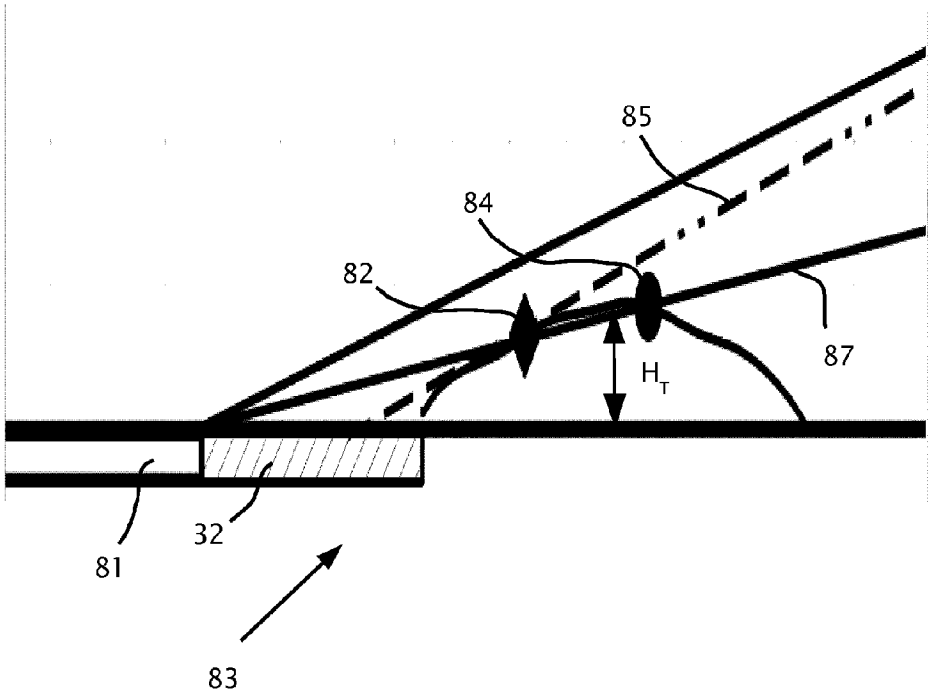


Fig. 8b

IPR2025-00943
Tesla EX1002 Page 1573

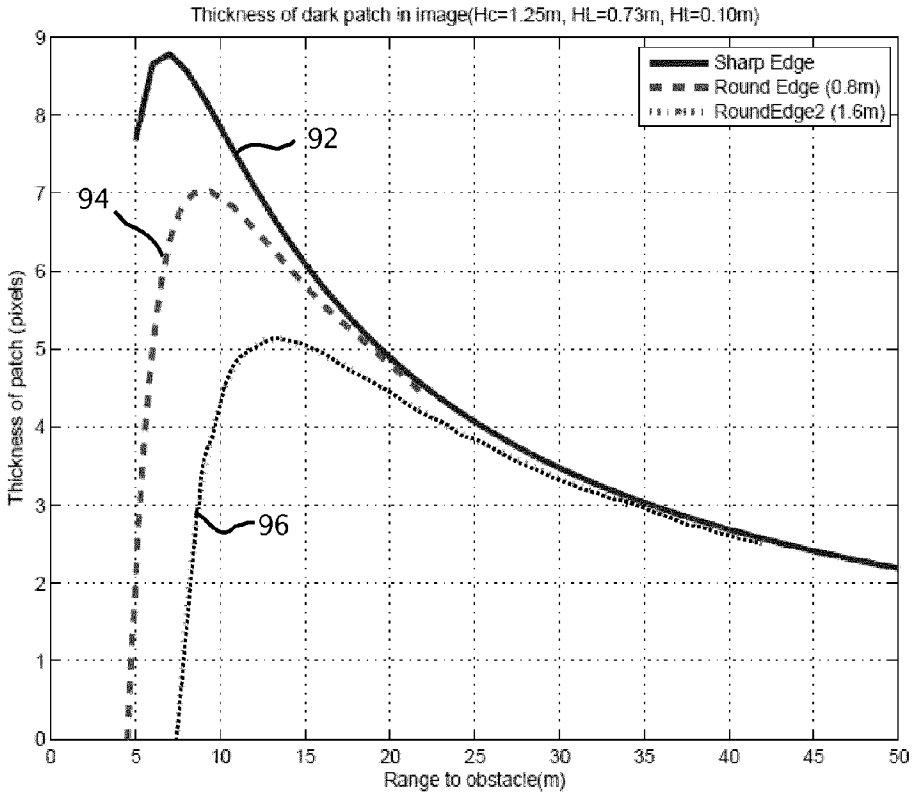


Fig. 9

IPR2025-00943
Tesla EX1002 Page 1574

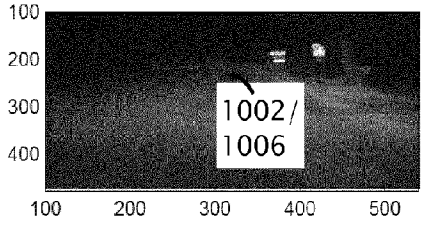


Fig. 10a

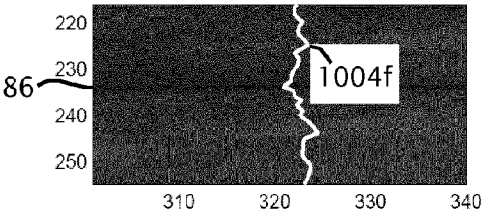


Fig. 10f

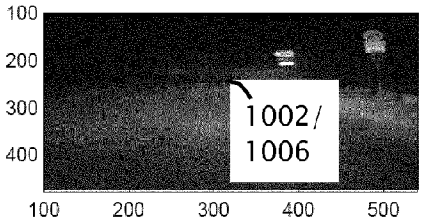


Fig. 10b

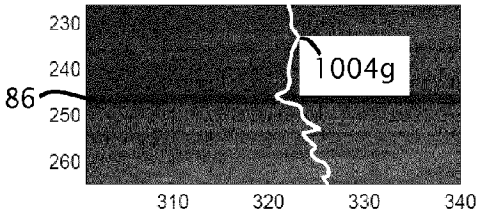


Fig. 10g

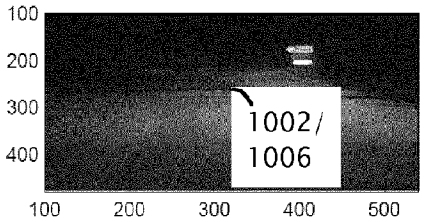


Fig. 10c

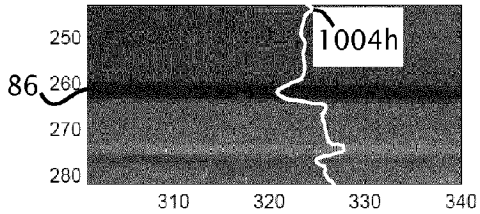


Fig. 10h

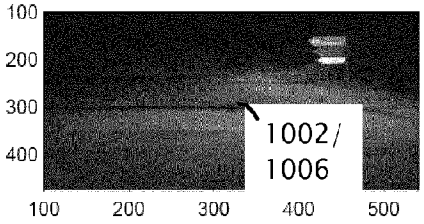


Fig. 10d

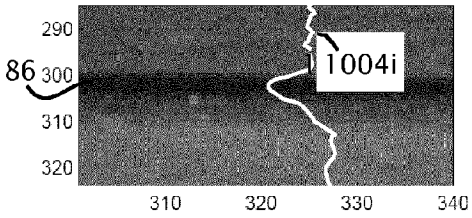


Fig. 10i

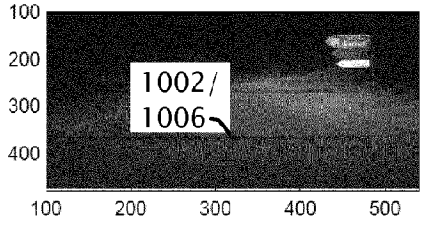


Fig. 10e

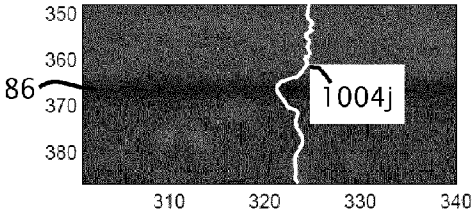


Fig. 10j

IPR2025-00943
Tesla EX1002 Page 1575

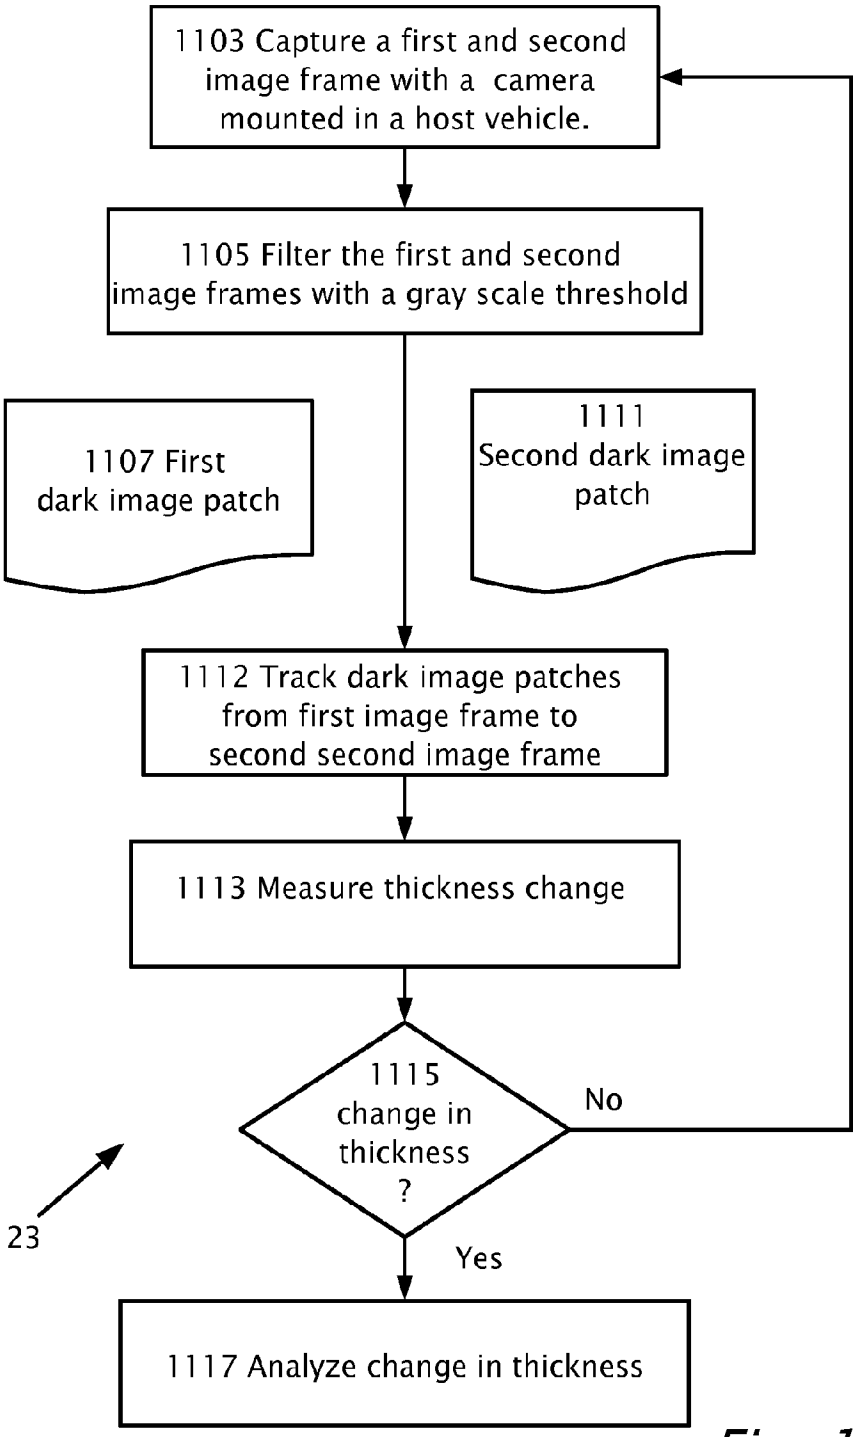


Fig. 11

IPR2025-00943
Tesla EX1002 Page 1576

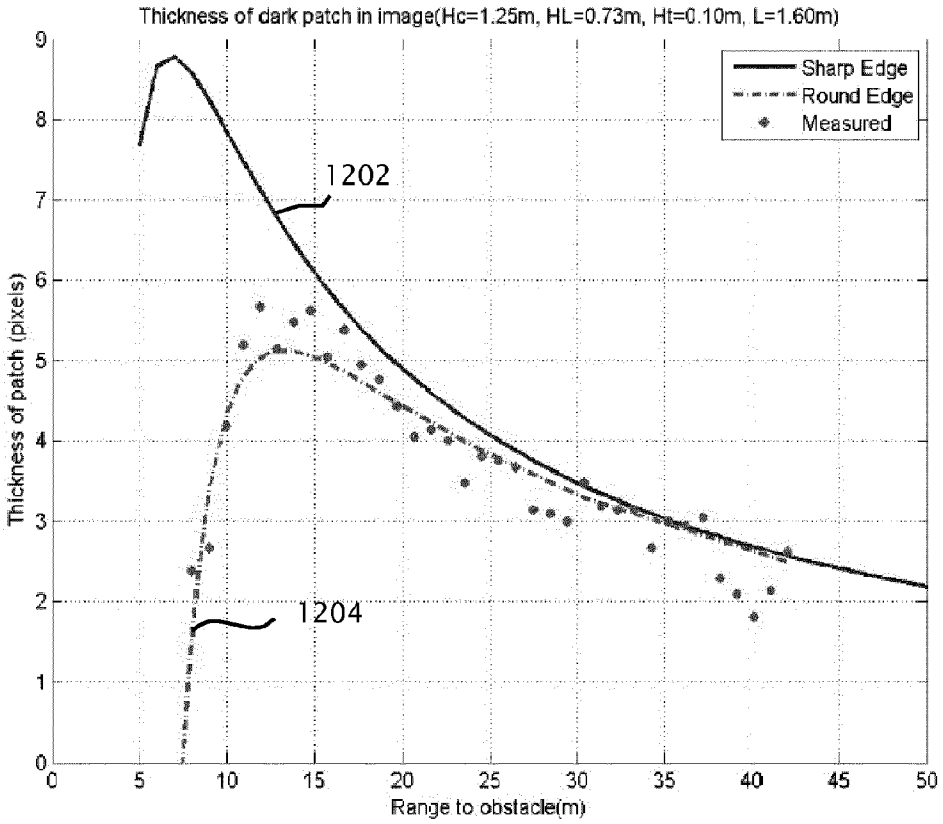


Fig. 12

IPR2025-00943
Tesla EX1002 Page 1577

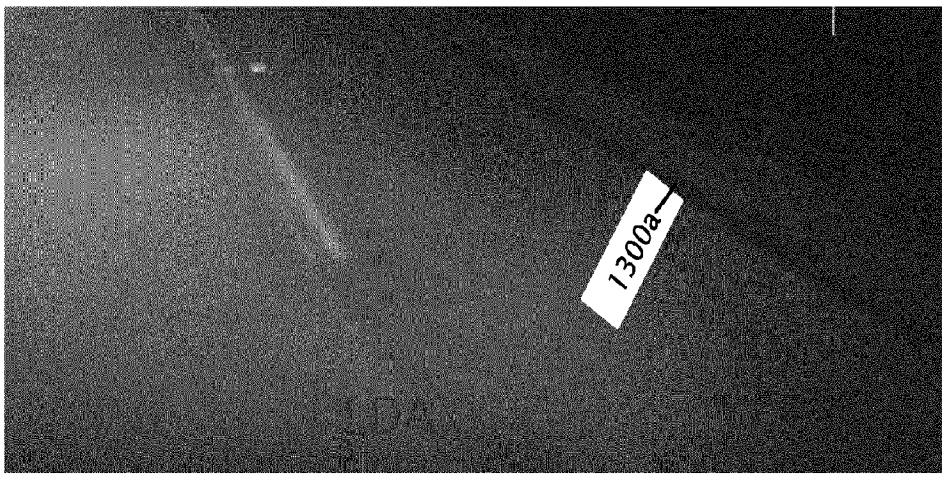


Fig. 13a

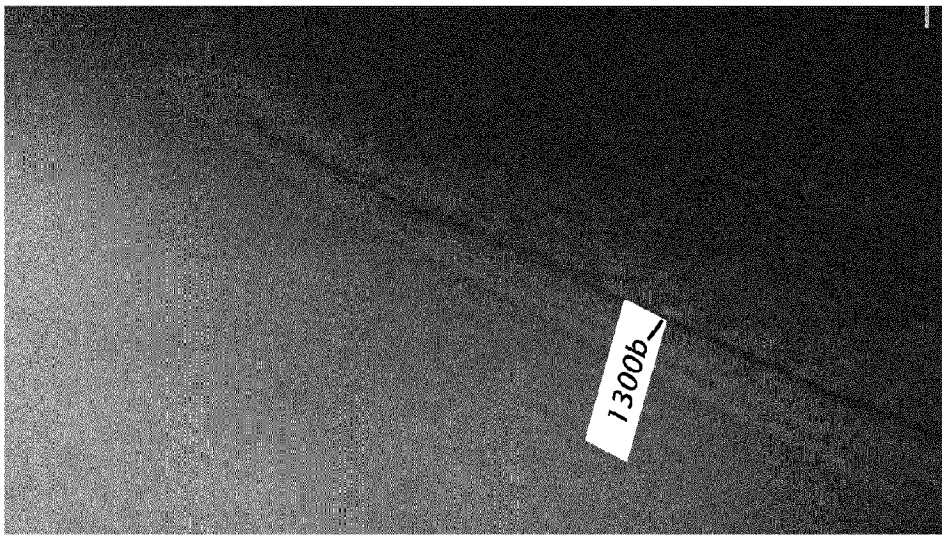


Fig. 13b

IPR2025-00943
Tesla EX1002 Page 1578

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 61569306 B **[0067]**
- US 61702755 B **[0067]**

Non-patent literature cited in the description

- **LARSON et al.** Lidar based off-road negative obstacle detection and analysis. *14th International IEEE Conference on Intelligent Transportation Systems,* October 2011 **[0003]**
- **BELUTTA et al.** Terrain Perception for Demo III. *Proceedings of the IEEE Intelligent Vehicles Symposium,* October 2000 **[0003]**
- Nighttime negative obstacle detection for off-road autonomous navigation In Unmanned Systems. **RABKIN et al.** Proceedings of the SPIE. 2007, vol. 6561, 656103 **[0003]**
- **MURAKA.** Detecting obstacles and drop-offs using stereo and motion cues for safe local motion. *IEEE/RSJ International Conference on Intelligent Robots and Systems,* September 2008 **[0003]**
- Preliminary investigation into the use of stereo illumination to enhance mobile robot terrain perception. **WITUS et al.** Proc. SPIE Vol. 4364, p. 290-301, Unmanned Ground Vehicle Technology III. vol. 4364, 290-301 **[0003]**

IPR2025-00943
Tesla EX1002 Page 1579

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 2 605 185 A3

(12) EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
09.07.2014 Bulletin 2014/28

(43) Date of publication A2:
19.06.2013 Bulletin 2013/25

(21) Application number: **12275201.7**

(22) Date of filing: **12.12.2012**

(51) Int Cl.:
G06K 9/00 *(2006.01)*

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR
Designated Extension States:
BA ME

(30) Priority: **12.12.2011 US 201161569306 P**
19.09.2012 US 201261702755 P

(71) Applicant: **Mobileye Technologies Limited**
24854 Nicosia (CY)

(72) Inventor: **Stein, Gideon P**
96263 Jerusalem (IL)

(74) Representative: **Jehan, Robert**
Williams Powell
Staple Court
11 Staple Inn Buildings
London, WC1V 7QH (GB)

(54) **Detection of obstacles at night by analysis of shadows**

(57) A driver assistance systems mountable in a host vehicle while the host vehicle is moving forward with headlights on for detection of obstacles based on shadows. The driver assistance system includes a camera (12) operatively connectible to a processor (30). A first image frame and a second image frame are captured of a road. A first dark image patch and a second dark image patch include intensity values less than a threshold value. The first and the second dark image patches are tracked from the first image frame to the second image frame as corresponding images of the same portion of the road. Respective thicknesses in vertical image coordinates are measured of the first and second dark image patches responsive to the tracking.

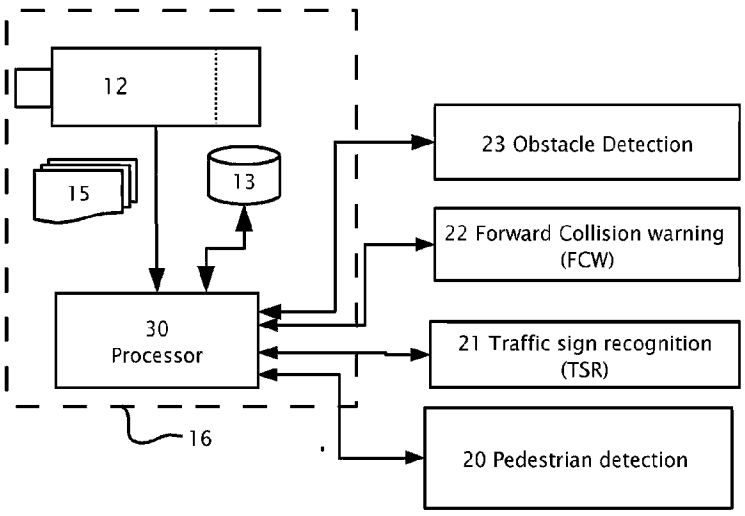


Fig. 2

EP 2 605 185 A3

Europäisches Patentamt
European Patent Office
Office européen des brevets

EUROPEAN SEARCH REPORT

Application Number
EP 12 27 5201

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	EP 1 708 125 A1 (VALEO VISION [FR]) 4 October 2006 (2006-10-04)	1-3,9-11	INV. G06K9/00
A	* paragraphs [0008], [0016] * * paragraphs [0018], [0021] - paragraph [0022] * * column 0024 * * paragraph [0026] - paragraph [0027] * * claim 18 *	4-8, 12-15	
A	HU TINGBO ET AL: "Negative obstacle detection from image sequences", THIRD INTERNATIONAL CONFERENCE ON DIGITAL IMAGE PROCESSING (ICDIP 2011), SPIE, 1000 20TH ST. BELLINGHAM WA 98225-6705 USA, vol. 8009, no. 1, 15 April 2011 (2011-04-15), pages 1-7, XP060015478, DOI: 10.1117/12.896288 [retrieved on 1901-01-01] * sections 2-3 *	1-15	

TECHNICAL FIELDS SEARCHED (IPC)
G06K

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
The Hague	30 May 2014	Kollreider, Klaus

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)

IPR2025-00943
Tesla EX1002 Page 1581

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 12 27 5201

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

30-05-2014

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 1708125 A1	04-10-2006	AT 425507 T	15-03-2009
		EP 1708125 A1	04-10-2006
		FR 2883818 A1	06-10-2006
		JP 2006298362 A	02-11-2006
		US 2006222208 A1	05-10-2006

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

IPR2025-00943
Tesla EX1002 Page 1582

Espacenet

Bibliographic data: EP3092599 (A1) — 2016-11-16

SYSTEMS AND METHODS FOR MIMICKING A LEADING VEHICLE

Inventor(s): CLARKE ANNA [IL]; BAGON EYAL [IL]; NIR GALI [IL] ± (CLARKE, ANNA, ; BAGON, EYAL, ; NIR, GALI, ; CLARKE, Anna, ; BAGON, Eyal, ; NIR, Gali)

Applicant(s): MOBILEYE VISION TECHNOLOGIES LTD [IL] ± (MOBILEYE VISION TECHNOLOGIES LTD, ; Mobileye Vision Technologies Ltd)

Classification: **- international:** ***B60K31/00; G06K9/00***

- cooperative: **B60K31/0066 (US)**; **B60R1/00 (US)**; **B60T13/662 (EP)**; **B60T7/12 (US)**; **B60T7/22 (EP, US)**; **B60T8/32 (US)**; **B60T8/321 (EP)**; **B60W10/18 (US)**; **B60W10/20 (US)**; **B60W30/00 (US)**; **B60W30/08 (US)**; **B60W30/09 (US)**; **B60W30/0953 (US)**; **B60W30/12 (US)**; **B60W30/14 (US)**; **B60W30/143 (US)**; **B60W30/165 (EP, US)**; **B60W30/18154 (EP, US)**; **B60W30/18163 (EP, US)**; **B60W40/06 (US)**; **B60W40/072 (US)**; **B60W40/076 (US)**; **B60W50/14 (US)**; **B62D15/025 (EP)**; **B62D15/0255 (EP)**; **B62D15/0265 (EP, US)**; **G01C11/04 (US)**; **G01C21/10 (US)**; **G01C21/26 (US)**; **G01C21/30 (EP, US)**; **G05D1/0088 (US)**; **G05D1/0246 (US)**; **G06K9/00798 (EP, US)**; **G06K9/00805 (EP)**; **G06K9/00825 (EP)**; **G06T7/285 (US)**; **G06T7/292 (US)**; **G06T7/70 (US)**; **G06T7/97 (US)**; **G08G1/09623 (EP, US)**; **G08G1/166 (EP, US)**; **G08G1/167 (EP, US)**; **H04N5/247 (US)**; B60K2031/0025 (US); B60T2201/022 (EP, US); B60T2201/03 (EP); B60T2201/08 (EP, US); B60T2210/34 (EP, US); B60W2050/0075 (EP, US); B60W2420/403 (US); B60W2420/42 (EP, US); B60W2510/0604 (US); B60W2510/18 (US); B60W2510/20 (US); B60W2520/14 (US); B60W2552/30 (EP); B60W2554/4041 (EP, US); B60W2554/80 (US); B60W2556/10 (EP, US); B60W2556/50 (EP, US); B60W2710/18 (US); B60W2710/20 (US); B62D15/0255 (US); B62D6/00 (US); G01S5/0027 (US); G05D1/0214 (US); G05D1/0253 (US); G05D1/0278 (US); G06T2207/10028 (US); G06T2207/30256 (US); G06T2207/30261 (US); G08G1/09 (US); G08G1/22 (EP, US); H04N7/18 (US)

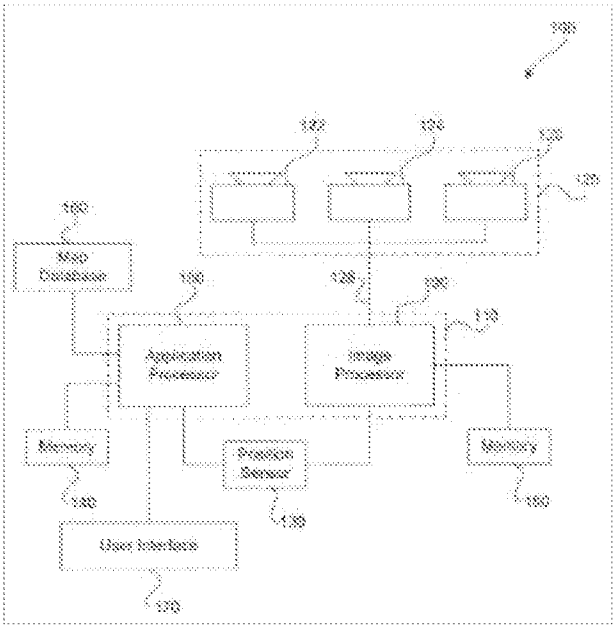
IPR2025-00943
Tesla EX1002 Page 1583

Application number: EP20140843184 20141204 Global Dossier

Priority number(s): US201361911490P 20131204 ; US201461993084P 20140514 ; US201461993111P 20140514 ; US201462015524P 20140623 ; US201462022221P 20140709 ; US201462040224P 20140821 ; US201462040269P 20140821 ; WO2014IB03151 20141204

Also published as: EP3092599 (B1) CN106030609 (A) CN106030609 (B) CN110920609 (A) CN111027420 (A) more

Abstract not available for EP3092599 (A1)
Abstract of corresponding document: US2015151725 (A1)



Systems and methods use cameras to provide autonomous navigation features. In one implementation, a driver assist navigation system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; identify, based on analysis of the plurality of images, a trigger for stopping the vehicle; and based on the identified trigger, cause the vehicle to stop according to a braking profile including a first segment associated with a first deceleration rate, a second segment which includes a second deceleration rate less than the first deceleration rate, and a third segment in which a level of braking is decreased as a target stopping location is approached, as determined based on the analysis of the plurality of images.

Espacenet

Description: EP3092599 (A1) — 2016-11-16

SYSTEMS AND METHODS FOR MIMICKING A LEADING VEHICLE

Description of EP3092599 (B1)

A high quality text as facsimile in your desired language may be available amongst the following family members:

CN106030609 (A) US2015151725 (A1) WO2015083009 (A1) CN110920609 (A) CN111027420 (A) US2015151742 (A1) US2015151751 (A1) US2015151753 (A1) US2015153735 (A1) US2015367850 (A1) US2016052514 (A1) US2016167650 (A1) US2016229410 (A1) US2016347322 (A1) US2019248372 (A1) US2020094839 (A1) US2020108834 (A1) US2020331477 (A1) more

The EPO does not accept any responsibility for the accuracy of data and information originating from other authorities than the EPO; in particular, the EPO does not guarantee that they are complete, up-to-date or fit for specific purposes.

BACKGROUND

I. Technical Field

[0001] The present disclosure relates generally to autonomous vehicle navigation and, more specifically, to systems and methods that use cameras to provide autonomous vehicle navigation features.

II. Background Information

[0002] As technology continues to advance, the goal of a fully autonomous vehicle that is capable of navigating on roadways is on the horizon. Primarily, an autonomous vehicle may be able to identify its environment and navigate without input from a human operator. Autonomous vehicles may also take into account a variety of factors and make appropriate decisions based on those factors to safely and accurately reach an intended destination. For example, various objects-such as other vehicles and pedestrians-are encountered when a vehicle typically travels a roadway. Autonomous driving systems may recognize these objects in a vehicle's environment and take appropriate and timely action to avoid collisions. Additionally, autonomous driving systems may identify other indicators-such as traffic signals, traffic signs, and lane markings-that regulate vehicle movement (e.g., when the vehicle must stop and may go, a speed at which the vehicle must not exceed, where the vehicle must be positioned on the roadway, etc.). Autonomous driving systems may need to determine

when a vehicle should change lanes, turn at intersections, change roadways, etc. As is evident from these examples, many factors may need to be addressed in order to provide an autonomous vehicle that is capable of navigating safely and accurately.

[0003] " Vision-based Long-distance Lane Perception and Front Vehicle Location for Full Autonomous Vehicles on Highway Roads" by Xin Liu , et al is directed to a vision-based long-distance lane perception and front location method for decision making by full autonomous vehicles on highway roads.

[0004] " A Transportable Neural-Network Approach to Autonomous Vehicle Following" by Nasser Kehtarnavaz , et al is directed to utilizing neural networks to control the speed and steering of a vehicle for the purpose of performing autonomous vehicle following, while following a leading vehicle.

[0005] US 2011/0190972 A1 is directed to systems for detecting the presence of stationary and non-stationary objects in the vicinity of a travelling vehicle, and controlling vehicle operational parameters in response to presence of such objects.

[0006] US 8,504,233 B1 is directed to a method for controlling a vehicle in an autonomous mode, which includes obtaining lane information that provides an estimated location of a lane on a road on which the vehicle is travelling and, if the lane information has become unavailable or unreliable, using sensor(s) to monitor neighboring vehicle(s) and controlling the vehicle to maintain at least a predetermined minimum distance between the vehicle and the neighboring vehicle(s).

[0007] JP2012-168592 is directed to a method for environment recognition of a vehicle. It comprises recognizing an arrow signal of a traffic light which is located far from the vehicle. It can also comprise performing a preceding vehicle detection and means to detect lanes detection or the road surface.

SUMMARY

[0008] Embodiments consistent with the present disclosure provide systems and methods for autonomous vehicle navigation. The disclosed embodiments may use cameras to provide autonomous vehicle navigation features. For example, consistent with the disclosed embodiments, the disclosed systems may include one, two, or more cameras that monitor the environment of a vehicle and cause a navigational response based on an analysis of images captured by one or more of the cameras.

[0009] In accordance with aspects of the invention there are provided a driver assist navigation system, a vehicle, and a method for navigating a vehicle as set out in claims 1, 10, and 11 respectively.

[0010] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the vehicle; determine from the plurality of images a second lane constraint on a second side of the vehicle opposite to the first side of the vehicle, wherein the first and second lane constraints define a lane within which the vehicle travels and wherein a first distance corresponds

Tesla EX1002 Page 1586 IPR2025-00943

to a distance between the first side of the vehicle and the first lane constraint and a second distance corresponds to a distance between the second side of the vehicle and the second lane constraint; determine, based on the plurality of images, whether a lane offset condition exists on the first side of the vehicle; if a lane offset condition exists on the first side of the vehicle, cause the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance; determine, based on the plurality of images, whether a lane offset condition exists on the second side of the vehicle; and if a lane offset condition exists on the second side of the vehicle, cause the vehicle to travel within the first and second lane constraints such that the first distance is less than the second distance.

[0011] Consistent with the disclosed embodiment, the at least one processing device may be further configured to cause the first distance to be substantially the same as the second distance if, based on the plurality of images, a lane offset condition is determined to exist on both the first and second sides of the vehicle. The at least one processing device may be configured to determine, based on the plurality of images, whether a lane offset condition is substantially non-existent, and if the lane offset condition is substantially non-existent, cause the vehicle to travel within the first and second lane constraints such that the first distance is substantially equal to the second distance. If the driver assist navigation system determines that a lane offset condition exists, the at least one processing device may be configured to cause the first distance to be at least 1.25 times greater than the second distance. Alternatively, the at least one processing device may be further configured to cause the first distance to be at least 2 times greater than the second distance. If the driver assist navigation system determines that the lane offset condition includes a curved road, the at least one processing device may be configured to cause the vehicle to travel within the first and second lane constraints such that the first distance is toward an outside of the curve relative to the vehicle, and the second distance is toward an inside of the curve relative to the vehicle. If the driver assist navigation system determines that the lane offset condition includes a presence of one or more objects on one side of the lane, the at least one processing device may be configured to cause the vehicle to travel within the first and second lane constraints such that the first distance, which is greater than the second distance, is provided toward the one or more objects. The one or more objects may include parked cars, a wall, and/or pedestrians, and may have a height of greater than 10 cm from a road surface. At least one of the first and second lane constraints may include a line marked on a road surface, an edge of a road surface, or a determined midpoint of a road surface width. Determination of whether a lane offset condition exists on the first side of the vehicle may be performed based on a position of a target object on the first side of the vehicle. The causing of the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance may be overridden based on input from a driver of the vehicle.

[0012] Consistent with another disclosed embodiment, a vehicle may include a first vehicle side; a second vehicle side opposite the first vehicle side; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on the first vehicle side; determine from the plurality of images a second lane constraint on the second vehicle side, wherein the first and second lane constraints define a lane within which the vehicle travels and wherein a first distance corresponds to a distance between the first vehicle side and the first lane constraint and a second distance

corresponds to a distance between the second vehicle side and the second lane constraint; determine, based on the plurality of images, whether a lane offset condition exists on the first vehicle side; if a lane offset condition exists on the first vehicle side, cause the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance; determine, based on the plurality of images, whether a lane offset condition exists on the second vehicle side; and if a lane offset condition exists on the second vehicle side, cause the vehicle to travel within the first and second lane constraints such that the first distance is less than the second distance.

[0013] Consistent with the disclosed embodiment, the at least one processing device may be further configured to cause the first distance to be substantially the same as the second distance if, based on the plurality of images, a lane offset condition is determined to exist on both the first and second vehicle sides. The at least one processing device may be configured to determine, based on the plurality of images, whether a lane offset condition is substantially non-existent, and if the lane offset condition is substantially non-existent, cause the vehicle to travel within the first and second lane constraints such that the first distance is substantially equal to the second distance. The lane offset condition may include a curved road, and the at least one processing device may be configured to cause the vehicle to travel within the first and second lane constraints such that the first distance is toward an outside of the curve relative to the vehicle, and the second distance is toward an inside of the curve relative to the vehicle. The lane offset condition may include a presence of one or more objects on one side of the lane, and the at least one processing device may be configured to cause the vehicle to travel within the first and second lane constraints such that the first distance, which is greater than the second distance, is provided toward the one or more objects. At least one of the first and second lane constraints may include an edge of a road surface and/or a determined midpoint of a road surface width. Determination of whether a lane offset condition exists on the first side of the vehicle may be performed based on a position of a target object on the first side of the vehicle. The causing of the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance may be overridden based on input from a driver of the vehicle.

[0014] Consistent with another disclosed embodiment, a method is provided for navigating a vehicle. The method may include acquiring, using at least one image capture device, a plurality of images of an area in the vicinity of the vehicle; determining from the plurality of images a first lane constraint on a first side of the vehicle; determining from the plurality of images a second lane constraint on a second side of the vehicle opposite to the first side of the vehicle, wherein the first and second lane constraints define a lane within which the vehicle travels and wherein a first distance corresponds to a distance between the first side of the vehicle and the first lane constraint and a second distance corresponds to a distance between the second side of the vehicle and the second lane constraint; determining, based on the plurality of images, whether a lane offset condition exists on the first side of the vehicle; if a lane offset condition exists on the first side of the vehicle, causing the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance; determining, based on the plurality of images, whether a lane offset condition exists on the second side of the vehicle; and if a lane offset condition exists on the second side of the vehicle, causing the vehicle to travel within the first and second lane constraints such that the first distance is less than the second distance.

[0015] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a current lane of travel from among a plurality of available travel lanes; and cause the vehicle to change lanes if the current lane of travel is not the same as a predetermined default travel lane.

[0016] Consistent with the disclosed embodiment, the predetermined default travel lane may be a right-most lane among the plurality of available travel lanes. The at least one processing device may be configured to activate a turn signal prior to causing the vehicle to change lanes. The at least one processing device may be configured to cause an audible announcement prior to causing the vehicle to change lanes. The at least one processing device may be configured to cause an audible indicator prior to causing the vehicle to change lanes. The predetermined default travel lane may be determined based on the plurality of images of the area in the vicinity of the vehicle, based on data received from a global positioning system, and/or based on an input received from a user.

[0017] Consistent with another disclosed embodiment, a vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a current lane of travel from among a plurality of available travel lanes; and cause the vehicle to change lanes if the current lane of travel is not the same as a predetermined default travel lane.

[0018] Consistent with the disclosed embodiment, the predetermined default travel lane may be a right-most lane among the plurality of available travel lanes. The at least one processing device may be configured to activate a turn signal prior to causing the vehicle to change lanes. The at least one processing device may be configured to cause an audible announcement prior to causing the vehicle to change lanes. The at least one processing device may be configured to cause an audible indicator prior to causing the vehicle to change lanes. The predetermined default travel lane may be determined based on the plurality of images of the area in the vicinity of the vehicle, based on data received from a global positioning system, and/or based on an input received from a user.

[0019] Consistent with another disclosed embodiment, a method is provided for navigating a vehicle. The method may include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the vehicle; determining from the plurality of images a current lane of travel from among a plurality of available travel lanes; and causing the vehicle to change lanes if the current lane of travel is not the same as a predetermined default travel lane.

[0020] Consistent with the disclosed embodiment, the predetermined default travel lane may be a right-most lane among the plurality of available travel lanes. A turn signal may be activated prior to causing the vehicle to change lanes. The predetermined default travel lane may be determined based on the plurality of images of the area in the vicinity of the vehicle.

Tesla EX1002 Page 1589 IPR2025-00943

IPR2025-00943
Tesla EX1002 Page 1590

[0021] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; recognize a curve to be navigated based on map data and vehicle position information; determine an initial target velocity for the vehicle based on at least one characteristic of the curve as reflected in the map data; adjust a velocity of the vehicle to the initial target velocity; determine, based on the plurality of images, one or more observed characteristics of the curve; determine an updated target velocity based on the one or more observed characteristics of the curve; and adjust the velocity of the vehicle to the updated target velocity.

[0022] Consistent with the disclosed embodiment, the at least one characteristic of the curve may include an expected radius of curvature for the curve as determined from the map data. The initial target velocity may be determined based on the expected radius of curvature for the curve and a lateral acceleration limit. The lateral acceleration limit may be between about 0.2 m/sec<2>and 3.0 m/sec<2>. Alternatively, the lateral acceleration limit may be between about 1.5 m/sec<2>and 3.0 m/sec<2>. The lateral acceleration limit may be adjusted based on a road width determined based on the plurality of images. The one or more observed characteristics of the curve may include an actual radius of curvature for the curve as determined through analysis of the plurality of images. The updated target velocity may be determined based on the actual radius of curvature for the curve and a lateral acceleration limit. The one or more observed characteristics of the curve may include a rate of change in an actual radius of curvature for the curve as determined through analysis of the plurality of images. The one or more observed characteristics of the curve may include a degree of banking associated with the curve as determined through analysis of the plurality of images. The one or more observed characteristics of the curve may include information, determined through analysis of the plurality of images, relating to the curve conveyed by one or more road signs. The initial target velocity may also be based on a focal range of the at least one image capture device. The initial target velocity may also be based on whether a traffic light is determined to be located along the curve. The initial target velocity may also be based on whether a traffic light is determined to be located within a predetermined distance after the curve.

[0023] Consistent with another disclosed embodiment, a vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; recognize a curve to be navigated based on map data and vehicle position information; determine an initial target velocity for the vehicle based on at least one characteristic of the curve as reflected in the map data; adjust a velocity of the vehicle to the initial target velocity; determine, based on the plurality of images, one or more observed characteristics of the curve; determine an updated target velocity based on the one or more observed characteristics of the curve; and adjust the velocity of the vehicle to the updated target velocity.

[0024] Consistent with the disclosed embodiment, the at least one characteristic of the curve may include an expected radius of curvature for the curve as determined from the map data. The initial target velocity may be determined based on the expected radius of curvature for the curve and a lateral acceleration limit. The lateral acceleration

limit may be adjusted based on a road width determined based on the plurality of images. The one or more observed characteristics of the curve may include an actual radius of curvature for the curve as determined through analysis of the plurality of images. The updated target velocity may be determined based on the actual radius of curvature for the curve and a lateral acceleration limit. The one or more observed characteristics of the curve may include a rate of change in an actual radius of curvature for the curve as determined through analysis of the plurality of images. The one or more observed characteristics of the curve may include a degree of banking associated with the curve as determined through analysis of the plurality of images. The one or more observed characteristics of the curve may include information, determined through analysis of the plurality of images, relating to the curve conveyed by one or more road signs. The initial target velocity may also be based on a focal range of the at least one image capture device. The initial target velocity may also be based on whether a traffic light is determined to be located along the curve. The initial target velocity may also be based on whether a traffic light is determined to be located within a predetermined distance after the curve.

[0025] Consistent with another disclosed embodiment, a method is provided for navigating a vehicle. The method may include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the vehicle; recognizing a curve to be navigated based on map data and vehicle position information; determining an initial target velocity for the vehicle based on at least one characteristic of the curve as reflected in the map data; adjusting a velocity of the vehicle to the initial target velocity; determining, based on the plurality of images, one or more observed characteristics of the curve; determining an updated target velocity based on the one or more observed characteristics of the curve; and adjusting the velocity of the vehicle to the updated target velocity.

[0026] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a primary vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the primary vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine, from at least some of the plurality of images, a first lane constraint on a first side of the primary vehicle; determine, from at least some of the plurality of images, a second lane constraint on a second side of the primary vehicle opposite to the first side of the primary vehicle, wherein the first and second lane constraints define a lane within which the primary vehicle travels; cause the primary vehicle to travel within the first and second lane constraints; locate in the plurality of images a leading vehicle; determine, based on the plurality of images, at least one action taken by the leading vehicle; and cause the primary vehicle to mimic the at least one action of the leading vehicle.

[0027] Consistent with the disclosed embodiment, at least one of the first and second lane constraints may include a line marked on a road surface. At least one of the first and second lane constraints may include an edge of a road surface. At least one of the first and second lane constraints may include a determined midpoint of a road surface width. The at least one action taken by the leading vehicle may include a turn, and the at least one processing device may cause the primary vehicle to execute a similar turn beginning at a location at which the leading vehicle initiated its turn. The at least one action taken by the leading vehicle may include an acceleration. The at least one action taken by the leading vehicle may include braking. The at least one action taken by the

leading vehicle may include providing a first offset distance on a side of the leading vehicle adjacent to the first lane constraint that is different from a second offset distance on a side of the leading vehicle adjacent to the second lane constraint. The at least one processing device may be further configured to determine, based on the plurality of images or based on map data and position information, whether the leading vehicle is turning at an intersection, and if the leading vehicle is determined to be turning at an intersection, decline to mimic the turn of the leading vehicle. The at least one processing device may be configured to determine, based on the plurality of images or based on map data and position information, whether the leading vehicle is turning at an intersection, determine whether the leading vehicle is changing lanes within the intersection, mimic the turn of the leading vehicle if the leading vehicle is determined to be turning at an intersection and not changing lanes within the intersection, and decline to mimic the turn of the leading vehicle if the leading vehicle is determined to be turning at an intersection, but changing lanes within the intersection. The at least one processing device may be configured to acquire a first leading vehicle based on the plurality of images, acquire a second leading vehicle based on the plurality of images, and cause the primary vehicle to mimic the first leading vehicle if, based on the plurality of images, the first leading vehicle is determined to be continuing on a path having less of a turning radius than a path followed by the second leading vehicle. The at least one processing device may be configured to determine position information for the leading vehicle and compare the position information to predetermined map data, create a navigation history for the leading vehicle by tracking movements of the leading vehicle relative to the predetermined map data, and determine whether to cause the primary vehicle to mimic the at least one action of the leading vehicle based on the navigation history of the leading vehicle. The at least one processing device may be configured to cause the primary vehicle to mimic a turn of the leading vehicle, even where no turn is detected based on the plurality of images, when the navigation history indicates that the leading vehicle has successfully navigated at least one prior turn.

[0028] Consistent with another disclosed embodiment, a primary vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the primary vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine, from at least some of the plurality of images, a first lane constraint on a first side of the primary vehicle; determine, from at least some of the plurality of images, a second lane constraint on a second side of the primary vehicle opposite to the first side of the primary vehicle, wherein the first and second lane constraints define a lane within which the primary vehicle travels; cause the primary vehicle to travel within the first and second lane constraints; locate in the plurality of images a leading vehicle; determine, based on the plurality of images, at least one action taken by the leading vehicle; and cause the primary vehicle to mimic the at least one action of the leading vehicle.

[0029] Consistent with the disclosed embodiment, at least one of the first and second lane constraints may include a line marked on a road surface. At least one of the first and second lane constraints may include an edge of a road surface. At least one of the first and second lane constraints may include a determined midpoint of a road surface width. The at least one action taken by the leading vehicle may include a turn, and the at least one processing device may cause the primary vehicle to execute a similar turn beginning at a location at which the leading vehicle initiated its turn. The at least one action taken by the leading vehicle may include an acceleration. The at least one action

Tesla EX1002 Page 1592 IPR2025-00943

taken by the leading vehicle may include braking. The at least one action taken by the leading vehicle may include providing a first offset distance on a side of the leading vehicle adjacent to the first lane constraint that is different from a second offset distance on a side of the leading vehicle adjacent to the second lane constraint. The at least one processing device may be further configured to determine, based on the plurality of images or based on map data and position information, whether the leading vehicle is turning at an intersection, and if the leading vehicle is determined to be turning at an intersection, decline to mimic the turn of the leading vehicle. The at least one processing device may be configured to determine, based on the plurality of images or based on map data and position information, whether the leading vehicle is turning at an intersection, determine whether the leading vehicle is changing lanes within the intersection, mimic the turn of the leading vehicle if the leading vehicle is determined to be turning at an intersection and not changing lanes within the intersection, and decline to mimic the turn of the leading vehicle if the leading vehicle is determined to be turning at an intersection, but changing lanes within the intersection. The at least one processing device may be configured to acquire a first leading vehicle based on the plurality of images, acquire a second leading vehicle based on the plurality of images, and cause the primary vehicle to mimic the first leading vehicle if, based on the plurality of images, the first leading vehicle is determined to be continuing on a path having less of a turning radius than a path followed by the second leading vehicle. The at least one processing device may be configured to determine position information for the leading vehicle and compare the position information to predetermined map data, create a navigation history for the leading vehicle by tracking movements of the leading vehicle relative to the predetermined map data, and determine whether to cause the primary vehicle to mimic the at least one action of the leading vehicle based on the navigation history of the leading vehicle. The at least one processing device may be configured to cause the primary vehicle to mimic a turn of the leading vehicle, even where no turn is detected based on the plurality of images, when the navigation history indicates that the leading vehicle has successfully navigated at least one prior turn.

[0030] Consistent with another disclosed embodiment, a method is provided for navigating a primary vehicle. The method ma include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the primary vehicle; determining, from at least some of the plurality of images, a first lane constraint on a first side of the primary vehicle; determining, from at least some of the plurality of images, a second lane constraint on a second side of the primary vehicle opposite to the first side of the primary vehicle, wherein the first and second lane constraints define a lane within which the primary vehicle travels; causing the primary vehicle to travel within the first and second lane constraints; locating in the plurality of images a leading vehicle; determining, based on the plurality of images, at least one action taken by the leading vehicle; and causing the primary vehicle to mimic the at least one action of the leading vehicle.

[0031] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a user vehicle. They system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the user vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the user vehicle; determine from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the

Tesla EX1002 Page 1593 IPR2025-00943

first and second lane constraints define a lane within which the user vehicle travels; acquire, based on the plurality of images, a target vehicle; enable the user vehicle to pass the target vehicle if the target vehicle is determined to be in a lane different from the lane in which the user vehicle is traveling; monitor a position of the target vehicle based on the plurality of images; and cause the user vehicle to abort the pass before completion of the pass, if the target vehicle is determined to be entering the lane in which the user vehicle is traveling.

[0032] Consistent with the disclosed embodiment, causing the user vehicle to abort the pass may include causing an application of brakes in the user vehicle. Causing the user vehicle to abort the pass may include causing an audible announcement. Causing the user vehicle to abort the pass may include causing the user vehicle to stop accelerating. Monitoring the position of the target vehicle may include estimating a speed of the target vehicle. At least one of the first and second lane constraints may include a line marked on a road surface, an edge of a road surface, and/or a determined midpoint of a road surface width.

[0033] Consistent with another disclosed embodiment, a user vehicle may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the user vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the user vehicle; determine from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; acquire, based on the plurality of images, a target vehicle; enable the user vehicle to pass the target vehicle if the target vehicle is determined to be in a lane different from the lane in which the user vehicle is traveling; monitor a position of the target vehicle based on the plurality of images; and cause the user vehicle to abort the pass before completion of the pass, if the target vehicle is determined to be entering the lane in which the user vehicle is traveling.

[0034] Consistent with the disclosed embodiment, causing the user vehicle to abort the pass may include causing an application of brakes in the user vehicle. Causing the user vehicle to abort the pass may include causing an audible announcement. Causing the user vehicle to abort the pass may include causing the user vehicle to stop accelerating. Monitoring the position of the target vehicle may include estimating a speed of the target vehicle. At least one of the first and second lane constraints may include a line marked on a road surface, an edge of a road surface, and/or a determined midpoint of a road surface width.

[0035] Consistent with another disclosed embodiment, a method is provided for navigating a user vehicle. The method may include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the user vehicle; determining from the plurality of images a first lane constraint on a first side of the user vehicle; determining from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; acquiring, based on the plurality of images, a target vehicle; enabling the user vehicle to pass the target vehicle if the target vehicle is determined to be in a lane different from the lane in which the user vehicle is traveling; monitoring a position of the target vehicle based on the plurality of images; and causing the user vehicle to abort the pass

before completion of the pass, if the target vehicle is determined to be entering the lane in which the user vehicle is traveling.

[0036] Consistent with the disclosed embodiment, causing the user vehicle to abort the pass may include causing an application of brakes in the user vehicle. Causing the user vehicle to abort the pass may include causing the user vehicle to stop accelerating. At least one of the first and second lane constraints may include a line marked on a road surface, an edge of a road surface, and/or a determined midpoint of a road surface width.

[0037] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a user vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the user vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the user vehicle; determine from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; determine, based on the plurality of images, whether an encroaching vehicle is approaching from a side of the user vehicle; and cause the user vehicle to maintain a current velocity and to travel within the first and second lane constraints such that a first distance, on the side of the user vehicle that the encroaching vehicle is approaching from, is greater than a second distance.

[0038] Consistent with the disclosed embodiment, the at least one processing device may be further configured to cause the user vehicle to take evasive action if analysis of the plurality of images indicates that the encroaching vehicle has crossed at least one collision threshold. The evasive action may include changing lanes, altering course, and/or braking. The at least one collision threshold may include a minimum predetermined distance to maintain between the user vehicle and the encroaching vehicle. The at least one collision threshold may include a time to collision. The at least one processing device may be configured to cause the user vehicle to return to a center of the lane, within which the user vehicle originally traveled, after completion of the evasive action. The at least one image capture device may have a field of view of at least 140 degrees. The at least one processing device may be configured to cause the user vehicle to resume a preset speed after braking in response to an encroaching vehicle. The at least one processing device may be configured to cause the user vehicle to return to a lane location within the first and second lane constraints such that the first distance between the user vehicle and the first lane constraint is substantially equal to the second distance between the user vehicle and the second lane constraint after altering course in response to an encroaching vehicle.

[0039] Consistent with another disclosed embodiment, a user vehicle may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the user vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the user vehicle; determine from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; determine, based on the plurality of images, whether an encroaching

IPR2025-00943
Tesla EX1002 Page 1595

vehicle is approaching from a side of the user vehicle; and cause the user vehicle to maintain a current velocity and to travel within the first and second lane constraints such that a first distance, on the side of the user vehicle that the encroaching vehicle is approaching from, is greater than a second distance.

[0040] Consistent with the disclosed embodiment, the at least one processing device may be further configured to cause the user vehicle to take evasive action if analysis of the plurality of images indicates that the encroaching vehicle has crossed at least one collision threshold. The evasive action may include changing lanes, altering course, and/or braking. The at least one collision threshold may include a minimum predetermined distance to maintain between the user vehicle and the encroaching vehicle. The at least one collision threshold may include a time to collision. The at least one processing device may be configured to cause the user vehicle to return to a center of the lane, within which the user vehicle originally traveled, after completion of the evasive action. The at least one image capture device may have a field of view of at least 140 degrees. The at least one processing device may be configured to cause the user vehicle to resume a preset speed after braking in response to an encroaching vehicle. The at least one processing device may be configured to cause the user vehicle to return to a lane location within the first and second lane constraints such that the first distance between the user vehicle and the first lane constraint is substantially equal to the second distance between the user vehicle and the second lane constraint after altering course in response to an encroaching vehicle.

[0041] Consistent with another disclosed embodiment, a method for navigating a user vehicle is provided. The method may include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the user vehicle; determining from the plurality of images a first lane constraint on a first side of the user vehicle; determining from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; determining, based on the plurality of images, whether an encroaching vehicle is approaching from a side of the user vehicle; and causing the user vehicle to maintain a current velocity and to travel within the first and second lane constraints such that a first distance, on the side of the user vehicle that the encroaching vehicle is approaching from, is greater than a second distance.

[0042] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a primary vehicle. They system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the primary vehicle; a data interface; and at least one processing device. The at least processing device may be configured to: receive the plurality of images via the data interface; identify a target object within the plurality of images; monitor, via the plurality of images, a motion of the target object and a distance between the primary vehicle and the target object; determine an indicator of an intercept time between the primary vehicle and the target object based on the monitored motion and the distance between the primary vehicle and the target object; and cause a response in the primary vehicle based on a comparison of the intercept time to a plurality of predetermined intercept thresholds.

[0043] Consistent with the disclosed em" odiment, the target object may be a target vehicle traveling in a lane different from the primary vehicle. The intercept time may be based on an indicator of speed of the tracked object, as determined based on the plurality of images. The intercept time may be based on an indicator of an acceleration

IPR2025-00943
Tesla EX1002 Page 1596

of the tracked object, as determined based on the plurality of images. The at least one processing device may be further configured to determine whether the target object is traveling in the same lane as the primary vehicle, and forego the response in the primary vehicle if the target object is determined to be traveling in a lane different from the primary vehicle. Determination of whether the target object is traveling in the same lane as the primary vehicle may be based on the plurality of images. Determination of whether the target object is traveling in the same lane as the primary vehicle may be based on at least one of predetermined map data or yaw rate of the primary vehicle. The response may be based on an averaged motion profile associated with the target object if the indicator of intercept time indicates that the target object is outside a first intercept threshold relative to the vehicle. The response may be based on actual motion of the target object if the indicator of intercept time indicates that the target object is between the first intercept threshold and a second intercept threshold relative to the vehicle. The response may include an emergency avoidance action if the indicator of intercept time indicates that the target object is within the second intercept threshold. The emergency avoidance action may include changing lanes and/or emergency braking.

[0044] Consistent with another disclosed embodiment, a primary vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the primary vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; identify a target object within the plurality of images; monitor, via the plurality of images, a motion of the target object and a distance between the primary vehicle and the target object; determine an indicator of an intercept time between the primary vehicle and the target object based on the monitored motion and the distance between the primary vehicle and the target object; and cause a response in the primary vehicle based on a comparison of the intercept time to a plurality of predetermined intercept thresholds.

[0045] Consistent with the disclosed embodiment, the target object may be a target vehicle traveling in a lane different from the primary vehicle. The intercept time may be based on an indicator of speed of the tracked object, as determined based on the plurality of images. The intercept time may be based on an indicator of an acceleration of the tracked object, as determined based on the plurality of images. The at least one processing device may be further configured to determine whether the target object is traveling in the same lane as the primary vehicle, and forego the response in the primary vehicle if the target object is determined to be traveling in a lane different from the primary vehicle. Determination of whether the target object is traveling in the same lane as the primary vehicle may be based on the plurality of images. Determination of whether the target object is traveling in the same lane as the primary vehicle may be based on at least one of predetermined map data or yaw rate of the primary vehicle. The response may be based on an averaged motion profile associated with the target object if the indicator of intercept time indicates that the target object is outside a first intercept threshold relative to the vehicle. The response may be based on actual motion of the target object if the indicator of intercept time indicates that the target object is between the first intercept threshold and a second intercept threshold relative to the vehicle. The response may include an emergency avoidance action if the indicator of intercept time indicates that the target object is within the second intercept threshold. The emergency avoidance action may include changing lanes and/or emergency braking.

IPR2025-00943
Tesla EX1002 Page 1597

[0046] Consistent with another disclosed embodiment, a method for navigating a primary vehicle may include acquiring, via at least one image capture device, a plurality of images of an area in a vicinity of the primary vehicle; identifying a target object within the plurality of images; monitoring, based on the plurality of images, a motion of the target object and a distance between the primary vehicle and the target object; determining an indicator of an intercept time between the primary vehicle and the target object based on the monitored motion and the distance between the primary vehicle and the target object; and causing a response in the primary vehicle based on a comparison of the intercept time to a plurality of predetermined intercept thresholds.

[0047] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; identify, based on analysis of the plurality of images, a trigger for stopping the vehicle; and based on the identified trigger, cause the vehicle to stop according to a braking profile including a first segment associated with a first deceleration rate, a second segment which includes a second deceleration rate less than the first deceleration rate, and a third segment in which a level of braking is decreased as a target stopping location is approached, as determined based on the analysis of the plurality of images.

[0048] Consistent with the disclosed embodiment, reaching of the target stopping location may be accompanied by a release in braking pressure. The trigger may include an object. The object may include another vehicle in front of the vehicle. The trigger may include a traffic light signal, a stop sign, and/or a speed limit sign. The at least one processing device may be further configured to provide an audible announcement identifying the trigger. The at least one processing device may be configured to cause an audible announcement indicating that the vehicle is braking.

[0049] Consistent with another disclosed embodiment, a vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; identify, based on analysis of the plurality of images, a trigger for stopping the vehicle; and based on the identified trigger, cause the vehicle to stop according to a braking profile including a first segment associated with a first deceleration rate, a second segment which includes a second deceleration rate less than the first deceleration rate, and a third segment in which a level of braking is decreased as a target stopping location is approached, as determined based on the analysis of the plurality of images.

[0050] Consistent with the disclosed embodiment, reaching of the target stopping location may be accompanied by a release in braking pressure. The trigger may include an object. The object may include another vehicle in front of the vehicle. The trigger may include a traffic light signal, a stop sign, and/or a speed limit sign. The at least one processing device may be further configured to provide an audible announcement identifying the trigger. The at least one processing device may be configured to cause an audible announcement indicating that the vehicle is braking.

[0051] Consistent with another disclosed embodiment, a method is provided for navigating a vehicle. The method may include acquiring, via at least one image capture

IPR2025-00943
Tesla EX1002 Page 1598

plurality of images of an area in a vicinity of the vehicle; identifying, based on analysis of the plurality of images, a trigger for stopping the vehicle; and based on the identified trigger, causing the vehicle to stop according to a braking profile including a first segment associated with a first deceleration rate, a second segment which includes a second deceleration rate less than the first deceleration rate, and a third segment in which a level of braking is decreased as a target stopping location is approached, as determined based on the analysis of the plurality of images.

[0052] Consistent with other disclosed embodiments, non-transitory computer-readable storage media may store program instructions, which are executed by at least one processing device and perform any of the methods described herein.

[0053] The foregoing general description and the following detailed description are exemplary and explanatory.

BRIEF DESCRIPTION OF THE DRAWINGS

[0054] The accompanying drawings, which are incorporated in and constitute a part of this disclosure, illustrate various disclosed embodiments. In the drawings:
FIG. 1 is a diagrammatic representation of an exemplary system consistent with the disclosed embodiments.
FIG. 2A is a diagrammatic side view representation of an exemplary vehicle including a system consistent with the disclosed embodiments.
FIG. 2B is a diagrammatic top view representation of the vehicle and system shown in FIG. 2A consistent with the disclosed embodiments.
FIG. 2C is a diagrammatic top view representation of another embodiment of a vehicle including a system consistent with the disclosed embodiments.
FIG. 2D is a diagrammatic top view representation of yet another embodiment of a vehicle including a system consistent with the disclosed embodiments.
FIG. 2E is a diagrammatic top view representation of yet another embodiment of a vehicle including a system consistent with the disclosed embodiments.
FIG. 2F is a diagrammatic representation of exemplary vehicle control systems consistent with the disclosed embodiments.
FIG. 3A is a diagrammatic representation of an interior of a vehicle including a rearview mirror and a user interface for a vehicle imaging system consistent with the disclosed embodiments.
FIG. 3B is an illustration of an example of a camera mount that is configured to be positioned behind a rearview mirror and against a vehicle windshield consistent with the disclosed embodiments.
FIG. 3C is an illustration of the camera mount shown in FIG. 3B from a different perspective consistent with the disclosed embodiments.
FIG. 3D is an illustration of an example of a camera mount that is configured to be positioned behind a rearview mirror and against a vehicle windshield consistent with the disclosed embodiments.
FIG. 4 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.
FIG. 5A is a flowchart showing an exemplary process for causing one or more navigational responses based on monocular image analysis consistent with disclosed embodiments.
FIG. 5B is a flowchart showing an exemplary process for detecting one or more vehicles and/or pedestrians in a set of images consistent with the disclosed embodiments.

IPR2025-00943
Tesla EX1002 Page 1599

IPR2025-00943
Tesla EX1002 Page 1600

FIG. 5C is a flowchart showing an exemplary process for detecting road marks and/or lane geometry information in a set of images consistent with the disclosed embodiments.
FIG. 5D is a flowchart showing an exemplary process for detecting traffic lights in a set of images consistent with the disclosed embodiments.
FIG. 5E is a flowchart showing an exemplary process for causing one or more navigational responses based on a vehicle path consistent with the disclosed embodiments.
FIG. 5F is a flowchart showing an exemplary process for determining whether a leading vehicle is changing lanes consistent with the disclosed embodiments.
FIG. 6 is a flowchart showing an exemplary process for causing one or more navigational responses based on stereo image analysis consistent with the disclosed embodiments.
FIG. 7 is a flowchart showing an exemplary process for causing one or more navigational responses based on an analysis of three sets of images consistent with the disclosed embodiments.
FIG. 8A is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint in which a lane offset condition exists on the first vehicle side consistent with the disclosed embodiments.
FIG. 8B is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint in which a lane offset condition exists on the second vehicle side consistent with the disclosed embodiments.
FIG. 8C is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint on a curved road consistent with the disclosed embodiments.
FIG. 8D is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint in which a lane offset condition exists on both the first vehicle side and on the second vehicle side consistent with the disclosed embodiments.
FIG. 9 is a diagrammatic representation of the memory of an exemplary navigation system consistent with the disclosed embodiments.
FIG. 10 is a flowchart of an exemplary method for navigating a vehicle consistent with the disclosed embodiments.
FIG 11 is a diagrammatic representation of an exemplary vehicle including a navigation system traveling within a default lane consistent with the disclosed embodiments.
FIG 12A is a diagrammatic representation of the memory of an exemplary navigation system consistent with the disclosed embodiments.
FIG 12B is a flowchart of an exemplary process for navigating a vehicle consistent with the disclosed embodiments.
FIGS. 13A and 13B are diagrammatic representations of an exemplary vehicle approaching, and navigating, a curve with one or more characteristics consistent with the disclosed embodiments.
FIG. 14 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.
FIG. 15 is a flow chart showing an exemplary process for controlling the velocity of a vehicle based on a detected curve and observed characteristics of the curve consistent with disclosed embodiments.
FIGS. 16A-16D are diagrammatic representations of a primary vehicle mimicking one or more actions of a leading vehicle consistent with the disclosed embodiments.
FIG. 17 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.
FIG. 18 is a flow chart showing an exemplary process for causing a primary vehicle to

mimic one or more actions of a leading vehicle consistent with disclosed embodiments.
FIG. 19 is a flow chart showing an exemplary process for causing a primary vehicle to decline to mimic a turn of a leading vehicle consistent with disclosed embodiments.
FIG. 20 is a flow chart showing an exemplary process for causing a primary vehicle to mimic or decline a turn of a leading vehicle consistent with disclosed embodiments.
FIG. 21 is a flow chart showing another exemplary process for causing a primary vehicle to mimic one or more actions of a leading vehicle consistent with disclosed embodiments.
FIG. 22 is a flow chart showing an exemplary process for causing a primary vehicle to mimic one or more actions of a leading vehicle based on a navigation history consistent with disclosed embodiments.
FIG. 23 is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint on a roadway along on which another vehicle is traveling consistent with the disclosed embodiments.
FIG. 24A is a diagrammatic representation of a memory storing instructions for navigating a vehicle among encroaching vehicles consistent with the disclosed embodiments.
FIG. 24B is a flowchart of an exemplary process for navigating a vehicle among encroaching vehicles consistent with the disclosed embodiments.
FIG. 25 is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint on a roadway along on which another vehicle is traveling consistent with the disclosed embodiments.
FIG. 26 is a diagrammatic representation of a memory storing instructions for detecting and responding to traffic laterally encroaching on a vehicle consistent with the disclosed embodiments.
FIG. 27 is a flowchart of an exemplary process for detecting and responding to traffic laterally encroaching on a vehicle consistent with the disclosed embodiments.
FIGS. 28A and 28B are diagrammatic representations of causing a response in a primary vehicle consistent with the disclosed embodiments.
FIG. 29 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.
FIG. 30 is a flow chart showing an exemplary process for causing a response in a primary vehicle consistent with disclosed embodiments.
FIG. 31 is a flow chart showing an exemplary process for declining to cause a response in a primary vehicle consistent with disclosed embodiments.
FIG. 32A is a diagrammatic representation of an exemplary vehicle in which an object is present in the vicinity of vehicle consistent with the disclosed embodiments.
FIG. 32B illustrates an exemplary braking profile, consistent with the disclosed embodiments.
FIG. 33 is a diagrammatic representation of the memory of an exemplary navigation system consistent with the disclosed embodiments.
FIG. 34 is a flowchart of an exemplary method for implementing a braking profile while navigating a vehicle consistent with the disclosed embodiments.

DETAILED DESCRIPTION

[0055] The following detailed description refers to the accompanying drawings. Wherever possible, the same reference numbers are used in the drawings and the following description to refer to the same or similar parts. While several illustrative embodiments are described herein, modifications, adaptations and other implementations are possible. For example, substitutions, additions or modifications may be made to the components illustrated in the drawings, and the illustrative

IPR2025-00943
Tesla EX1002 Page 1601

methods described herein may be modified by substituting, reordering, removing, or adding steps to the disclosed methods. Accordingly, the following detailed description is not limited to the disclosed embodiments and examples. Instead, the proper scope is defined by the appended claims.

[0056] Disclosed embodiments provide systems and methods that use cameras to provide autonomous navigation features. In various embodiments, the system may include one, two or more cameras that monitor the environment of a vehicle. In one embodiment, the system may monitor and adjust the free space between a vehicle and the boundaries of the lane within which the vehicle is traveling. In another embodiment, the system may select a particular lane as a default lane for the vehicle to use while traveling. In another embodiment, the system may control the speed of the vehicle in different scenarios, such as while making a turn. In yet another embodiment, the system may mimic the actions of a leading vehicle. In yet another embodiment, the system may monitor a target vehicle and enable the vehicle to pass the target vehicle under certain conditions (e.g., if the target vehicle is traveling in a lane different from the lane within which the vehicle is traveling. In still yet another embodiment, the system may produce a natural feeling response to a laterally encroaching vehicle, such as a vehicle attempting to move into the lane within which the vehicle is traveling.

[0057] FIG. 1 is a block diagram representation of a system 100 consistent with the exemplary disclosed embodiments. System 100 may include various components depending on the requirements of a particular implementation. In some embodiments, system 100 may include a processing unit 110, an image acquisition unit 120, a position sensor 130, one or more memory units 140, 150, a map database 160, and a user interface 170. Processing unit 110 may include one or more processing devices. In some embodiments, processing unit 110 may include an applications processor 180, an image processor 190, or any other suitable processing device. Similarly, image acquisition unit 120 may include any number of image acquisition devices and components depending on the requirements of a particular application. In some embodiments, image acquisition unit 120 may include one or more image capture devices (e.g., cameras), such as image capture device 122, image capture device 124, and image capture device 126. System 100 may also include a data interface 128 communicatively connecting processing device 110 to image acquisition device 120. For example, data interface 128 may include any wired and/or wireless link or links for transmitting image data acquired by image accusation device 120 to processing unit 110.

[0058] Both applications processor 180 and image processor 190 may include various types of processing devices. For example, either or both of applications processor 180 and image processor 190 may include a microprocessor, preprocessors (such as an image preprocessor), graphics processors, a central processing unit (CPU), support circuits, digital signal processors, integrated circuits, memory, or any other types of devices suitable for running applications and for image processing and analysis. In some embodiments, applications processor 180 and/or image processor 190 may include any type of single or multi-core processor, mobile device microcontroller, central processing unit, etc. Various processing devices may be used, including, for example, processors available from manufacturers such as Intel®, AMD®, etc. and may include various architectures (e.g., x86 processor, ARM®, etc.).

[0059] In some embodiments, applications processor 180 and/or image processor 190 may include any of the EyeQ series of processor chips available from Mobileye®.

IPR2025-00943
Tesla EX1002 Page 1602

These processor designs each include multiple processing units with local memory and instruction sets. Such processors may include video inputs for receiving image data from multiple image sensors and may also include video out capabilities. In one example, the EyeQ2® uses 90nm-micron technology operating at 332Mhz. The EyeQ2® architecture consists of two floating point, hyper-thread 32-bit RISC CPUs (MIPS32® 34K® cores), five Vision Computing Engines (VCE), three Vector Microcode Processors (VMP®), Denali 64-bit Mobile DDR Controller, 128-bit internal Sonics Interconnect, dual 16-bit Video input and 18-bit Video output controllers, 16 channels DMA and several peripherals. The MIPS34K CPU manages the five VCEs, three VMP™ and the DMA, the second MIPS34K CPU and the multi-channel DMA as well as the other peripherals. The five VCEs, three VMP® and the MIPS34K CPU can perform intensive vision computations required by multi-function bundle applications. In another example, the EyeQ3®, which is a third generation processor and is six times more powerful that the EyeQ2®, may be used in the disclosed embodiments.

[0060] While FIG. 1 depicts two separate processing devices included in processing unit 110, more or fewer processing devices may be used. For example, in some embodiments, a single processing device may be used to accomplish the tasks of applications processor 180 and image processor 190. In other embodiments, these tasks may be performed by more than two processing devices.

[0061] Processing unit 110 may comprise various types of devices. For example, processing unit 110 may include various devices, such as a controller, an image preprocessor, a central processing unit (CPU), support circuits, digital signal processors, integrated circuits, memory, or any other types of devices for image processing and analysis. The image preprocessor may include a video processor for capturing, digitizing and processing the imagery from the image sensors. The CPU may comprise any number of microcontrollers or microprocessors. The support circuits may be any number of circuits generally well known in the art, including cache, power supply, clock and input-output circuits. The memory may store software that, when executed by the processor, controls the operation of the system. The memory may include databases and image processing software. The memory may comprise any number of random access memories, read only memories, flash memories, disk drives, optical storage, tape storage, removable storage and other types of storage. In one instance, the memory may be separate from the processing unit 110. In another instance, the memory may be integrated into the processing unit 110.

[0062] Each memory 140, 150 may include software instructions that when executed by a processor (e.g., applications processor 180 and/or image processor 190), may control operation of various aspects of system 100. These memory units may include various databases and image processing software. The memory units may include random access memory, read only memory, flash memory, disk drives, optical storage, tape storage, removable storage and/or any other types of storage. In some embodiments, memory units 140, 150 may be separate from the applications processor 180 and/or image processor 190. In other embodiments, these memory units may be integrated into applications processor 180 and/or image processor 190.

[0063] Position sensor 130 may include any type of device suitable for determining a location associated with at least one component of system 100. In some embodiments, position sensor 130 may include a GPS receiver. Such receivers can determine a user position and velocity by processing signals broadcasted by global positioning system satellites. Position information from position sensor 130 may be made available to

IPR2025-00943
Tesla EX1002 Page 1603

applications processor 180 and/or image processor 190.

[0064] User interface 170 may include any device suitable for providing information to or for receiving inputs from one or more users of system 100. In some embodiments, user interface 170 may include user input devices, including, for example, a touchscreen, microphone, keyboard, pointer devices, track wheels, cameras, knobs, buttons, etc. With such input devices, a user may be able to provide information inputs or commands to system 100 by typing instructions or information, providing voice commands, selecting menu options on a screen using buttons, pointers, or eye-tracking capabilities, or through any other suitable techniques for communicating information to system 100.

[0065] User interface 170 may be equipped with one or more processing devices configured to provide and receive information to or from a user and process that information for use by, for example, applications processor 180. In some embodiments, such processing devices may execute instructions for recognizing and tracking eye movements, receiving and interpreting voice commands, recognizing and interpreting touches and/or gestures made on a touchscreen, responding to keyboard entries or menu selections, etc. In some embodiments, user interface 170 may include a display, speaker, tactile device, and/or any other devices for providing output information to a user.

[0066] Map database 160 may include any type of database for storing map data useful to system 100. In some embodiments, map database 160 may include data relating to the position, in a reference coordinate system, of various items, including roads, water features, geographic features, businesses, points of interest, restaurants, gas stations, etc. Map database 160 may store not only the locations of such items, but also descriptors relating to those items, including, for example, names associated with any of the stored features. In some embodiments, map database 160 may be physically located with other components of system 100. Alternatively or additionally, map database 160 or a portion thereof may be located remotely with respect to other components of system 100 (e.g., processing unit 110). In such embodiments, information from map database 160 may be downloaded over a wired or wireless data connection to a network (e.g., over a cellular network and/or the Internet, etc.).

[0067] Image capture devices 122, 124, and 126 may each include any type of device suitable for capturing at least one image from an environment. Moreover, any number of image capture devices may be used to acquire images for input to the image processor. Some embodiments may include only a single image capture device, while other embodiments may include two, three, or even four or more image capture devices. Image capture devices 122,124, and 126 will be further described with reference to FIGS. 2B-2E , below.

[0068] System 100, or various components thereof, may be incorporated into various different platforms. In some embodiments, system 100 may be included on a vehicle 200, as shown in FIG. 2A . For example, vehicle 200 may be equipped with a processing unit 110 and any of the other components of system 100, as described above relative to FIG. 1 . While in some embodiments vehicle 200 may be equipped with only a single image capture device (e.g., camera), in other embodiments, such as those discussed in connection with FIGS. 2B-2E , multiple image capture devices may be used. For example, either of image capture devices 122 and 124 of vehicle 200, as shown in FIG. 2A , may be part of an ADAS (Advanced Driver Assistance Systems)

imaging set.

[0069] The image capture devices included on vehicle 200 as part of the image acquisition unit 120 may be positioned at any suitable location. In some embodiments, as shown in FIGS. 2A-2E and 3A-3C , image capture device 122 may be located in the vicinity of the rearview mirror. This position may provide a line of sight similar to that of the driver of vehicle 200, which may aid in determining what is and is not visible to the driver. Image capture device 122 may be positioned at any location near the rearview mirror, but placing image capture device 122 on the driver side of the mirror may further aid in obtaining images representative of the driver's field of view and/or line of sight.

[0070] Other locations for the image capture devices of image acquisition unit 120 may also be used. For example, image capture device 124 may be located on or in a bumper of vehicle 200. Such a location may be especially suitable for image capture devices having a wide field of view. The line of sight of bumper-located image capture devices can be different from that of the driver and, therefore, the bumper image capture device and driver may not always see the same objects. The image capture devices (e.g., image capture devices 122, 124, and 126) may also be located in other locations. For example, the image capture devices may be located on or in one or both of the side mirrors of vehicle 200, on the roof of vehicle 200, on the hood of vehicle 200, on the trunk of vehicle 200, on the sides of vehicle 200, mounted on, positioned behind, or positioned in front of any of the windows of vehicle 200, and mounted in or near light figures on the front and/or back of vehicle 200, etc.

[0071] In addition to image capture devices, vehicle 200 may include various other components of system 100. For example, processing unit 110 may be included on vehicle 200 either integrated with or separate from an engine control unit (ECU) of the vehicle. Vehicle 200 may also be equipped with a position sensor 130, such as a GPS receiver and may also include a map database 160 and memory units 140 and 150.

[0072] FIG. 2A is a diagrammatic side view representation of an exemplary vehicle imaging system consistent with the disclosed embodiments. FIG. 2B is a diagrammatic top view illustration of the embodiment shown in FIG. 2A . As illustrated in FIG. 2B , the disclosed embodiments may include a vehicle 200 including in its body a system 100 with a first image capture device 122 positioned in the vicinity of the rearview mirror and/or near the driver of vehicle 200, a second image capture device 124 positioned on or in a bumper region (e.g., one of bumper regions 210) of vehicle 200, and a processing unit 110.

[0073] As illustrated in FIG. 2C , image capture devices 122 and 124 may both be positioned in the vicinity of the rearview mirror and/or near the driver of vehicle 200. Additionally, while two image capture devices 122 and 124 are shown in FIGS. 2B and 2C , it should be understood that other embodiments may include more than two image capture devices. For example, in the embodiments shown in FIGS. 2D and 2E , first, second, and third image capture devices 122, 124, and 126, are included in the system 100 of vehicle 200.

[0074] As illustrated in FIG. 2D , image capture device 122 may be positioned in the vicinity of the rearview mirror and/or near the driver of vehicle 200, and image capture devices 124 and 126 may be positioned on or in a bumper region (e.g., one of bumper regions 210) of vehicle 200. And as shown in FIG. 2E , image capture devices 122, 124, and 126 may be positioned in the vicinity of the rearview mirror and/or near the

IPR2025-00943
Tesla EX1002 Page 1605

driver seat of vehicle 200. The disclosed embodiments are not limited to any particular number and configuration of the image capture devices, and the image capture devices may be positioned in any appropriate location within and/or on vehicle 200.

[0075] It is to be understood that the disclosed embodiments are not limited to vehicles and could be applied in other contexts. It is also to be understood that disclosed embodiments are not limited to a particular type of vehicle 200 and may be applicable to all types of vehicles including automobiles, trucks, trailers, and other types of vehicles.

[0076] The first image capture device 122 may include any suitable type of image capture device. Image capture device 122 may include an optical axis. In one instance, the image capture device 122 may include an Aptina M9V024 WVGA sensor with a global shutter. In other embodiments, image capture device 122 may provide a resolution of 1280x960 pixels and may include a rolling shutter. Image capture device 122 may include various optical elements. In some embodiments one or more lenses may be included, for example, to provide a desired focal length and field of view for the image capture device. In some embodiments, image capture device 122 may be associated with a 6mm lens or a 12mm lens. In some embodiments, image capture device 122 may be configured to capture images having a desired field-of-view (FOV) 202, as illustrated in FIG. 2D . For example, image capture device 122 may be configured to have a regular FOV, such as within a range of 40 degrees to 56 degrees, including a 46 degree FOV, 50 degree FOV, 52 degree FOV, or greater. Alternatively, image capture device 122 may be configured to have a narrow FOV in the range of 23 to 40 degrees, such as a 28 degree FOV or 36 degree FOV. In addition, image capture device 122 may be configured to have a wide FOV in the range of 100 to 180 degrees. In some embodiments, image capture device 122 may include a wide angle bumper camera or one with up to a 180 degree FOV.

[0077] The first image capture device 122 may acquire a plurality of first images relative to a scene associated with the vehicle 200. Each of the plurality of first images may be acquired as a series of image scan lines, which may be captured using a rolling shutter. Each scan line may include a plurality of pixels.

[0078] The first image capture device 122 may have a scan rate associated with acquisition of each of the first series of image scan lines. The scan rate may refer to a rate at which an image sensor can acquire image data associated with each pixel included in a particular scan line.

[0079] Image capture devices 122, 124, and 126 may contain any suitable type and number of image sensors, including CCD sensors or CMOS sensors, for example. In one embodiment, a CMOS image sensor may be employed along with a rolling shutter, such that each pixel in a row is read one at a time, and scanning of the rows proceeds on a row-by-row basis until an entire image frame has been captured. In some embodiments, the rows may be captured sequentially from top to bottom relative to the frame.

[0080] The use of a rolling shutter may result in pixels in different rows being exposed and captured at different times, which may cause skew and other image artifacts in the captured image frame. On the other hand, when the image capture device 122 is configured to operate with a global or synchronous shutter, all of the pixels may be exposed for the same amount of time and during a common exposure period. As a

Tesla EX1002 Page 1606 IPR2025-00943

result, the image data in a frame collected from a system employing a global shutter represents a snapshot of the entire FOV (such as FOV 202) at a particular time. In contrast, in a rolling shutter application, each row in a frame is exposed and data is capture at different times. Thus, moving objects may appear distorted in an image capture device having a rolling shutter. This phenomenon will be described in greater detail below.

[0081] The second image capture device 124 and the third image capturing device 126 may be any type of image capture device. Like the first image capture device 122, each of image capture devices 124 and 126 may include an optical axis. In one embodiment, each of image capture devices 124 and 126 may include an Aptina M9V024 WVGA sensor with a global shutter. Alternatively, each of image capture devices 124 and 126 may include a rolling shutter. Like image capture device 122, image capture devices 124 and 126 may be configured to include various lenses and optical elements. In some embodiments, lenses associated with image capture devices 124 and 126 may provide FOVs (such as FOVs 204 and 206) that are the same as, or narrower than, a FOV (such as FOV 202) associated with image capture device 122. For example, image capture devices 124 and 126 may have FOVs of 40 degrees, 30 degrees, 26 degrees, 23 degrees, 20 degrees, or less.

[0082] Image capture devices 124 and 126 may acquire a plurality of second and third images relative to a scene associated with the vehicle 200. Each of the plurality of second and third images may be acquired as a second and third series of image scan lines, which may be captured using a rolling shutter. Each scan line or row may have a plurality of pixels. Image capture devices 124 and 126 may have second and third scan rates associated with acquisition of each of image scan lines included in the second and third series.

[0083] Each image capture device 122, 124, and 126 may be positioned at any suitable position and orientation relative to vehicle 200. The relative positioning of the image capture devices 122, 124, and 126 may be selected to aid in fusing together the information acquired from the image capture devices. For example, in some embodiments, a FOV (such as FOV 204) associated with image capture device 124 may overlap partially or fully with a FOV (such as FOV 202) associated with image capture device 122 and a FOV (such as FOV 206) associated with image capture device 126.

[0084] Image capture devices 122, 124, and 126 may be located on vehicle 200 at any suitable relative heights. In one instance, there may be a height difference between the image capture devices 122, 124, and 126, which may provide sufficient parallax information to enable stereo analysis. For example, as shown in FIG. 2A , the two image capture devices 122 and 124 are at different heights. There may also be a lateral displacement difference between image capture devices 122, 124, and 126, giving additional parallax information for stereo analysis by processing unit 110, for example. The difference in the lateral displacement may be denoted by dx, as shown in FIGS. 2C and 2D . In some embodiments, fore or aft displacement (e.g., range displacement) may exist between image capture devices 122, 124, and 126. For example, image capture device 122 may be located 0.5 to 2 meters or more behind image capture device 124 and/or image capture device 126. This type of displacement may enable one of the image capture devices to cover potential blind spots of the other image capture device(s).

Tesla EX1002 Page 1607 IPR2025-00943

[0085] Image capture devices 122 may have any suitable resolution capability (e.g., number of pixels associated with the image sensor), and the resolution of the image sensor(s) associated with the image capture device 122 may be higher, lower, or the same as the resolution of the image sensor(s) associated with image capture devices 124 and 126. In some embodiments, the image sensor(s) associated with image capture device 122 and/or image capture devices 124 and 126 may have a resolution of 640 x 480, 1024 x 768, 1280 x 960, or any other suitable resolution.

[0086] The frame rate (e.g., the rate at which an image capture device acquires a set of pixel data of one image frame before moving on to capture pixel data associated with the next image frame) may be controllable. The frame rate associated with image capture device 122 may be higher, lower, or the same as the frame rate associated with image capture devices 124 and 126. The frame rate associated with image capture devices 122, 124, and 126 may depend on a variety of factors that may affect the timing of the frame rate. For example, one or more of image capture devices 122, 124, and 126 may include a selectable pixel delay period imposed before or after acquisition of image data associated with one or more pixels of an image sensor in image capture device 122, 124, and/or 126. Generally, image data corresponding to each pixel may be acquired according to a clock rate for the device (e.g., one pixel per clock cycle). Additionally, in embodiments including a rolling shutter, one or more of image capture devices 122, 124, and 126 may include a selectable horizontal blanking period imposed before or after acquisition of image data associated with a row of pixels of an image sensor in image capture device 122, 124, and/or 126. Further, one or more of image capture devices 122, 124, and/or 126 may include a selectable vertical blanking period imposed before or after acquisition of image data associated with an image frame of image capture device 122, 124, and 126.

[0087] These timing controls may enable synchronization of frame rates associated with image capture devices 122,124, and 126, even where the line scan rates of each are different. Additionally, as will be discussed in greater detail below, these selectable timing controls, among other factors (e.g., image sensor resolution, maximum line scan rates, etc.) may enable synchronization of image capture from an area where the FOV of image capture device 122 overlaps with one or more FOVs of image capture devices 124 and 126, even where the field of view of image capture device 122 is different from the FOVs of image capture devices 124 and 126.

[0088] Frame rate timing in image capture device 122,124, and 126 may depend on the resolution of the associated image sensors. For example, assuming similar line scan rates for both devices, if one device includes an image sensor having a resolution of 640 x 480 and another device includes an image sensor with a resolution of 1280 x 960, then more time will be required to acquire a frame of image data from the sensor having the higher resolution.

[0089] Another factor that may affect the timing of image data acquisition in image capture devices 122, 124, and 126 is the maximum line scan rate. For example, acquisition of a row of image data from an image sensor included in image capture device 122, 124, and 126 will require some minimum amount of time. Assuming no pixel delay periods are added, this minimum amount of time for acquisition of a row of image data will be related to the maximum line scan rate for a particular device. Devices that offer higher maximum line scan rates have the potential to provide higher frame rates than devices with lower maximum line scan rates. In some embodiments, one or more of image capture devices 124 and 126 may have a maximum line scan

IPR2025-00943
Tesla EX1002 Page 1608

rate that is higher than a maximum line scan rate associated with image capture device 122. In some embodiments, the maximum line scan rate of image capture device 124 and/or 126 may be 1.25, 1.5, 1.75, or 2 times or more than a maximum line scan rate of image capture device 122.

[0090] In another embodiment, image capture devices 122, 124, and 126 may have the same maximum line scan rate, but image capture device 122 may be operated at a scan rate less than or equal to its maximum scan rate. The system may be configured such that one or more of image capture devices 124 and 126 operate at a line scan rate that is equal to the line scan rate of image capture device 122. In other instances, the system may be configured such that the line scan rate of image capture device 124 and/or image capture device 126 may be 1.25, 1.5, 1.75, or 2 times or more than the line scan rate of image capture device 122.

[0091] In some embodiments, image capture devices 122, 124, and 126 may be asymmetric. That is, they may include cameras having different fields of view (FOV) and focal lengths. The fields of view of image capture devices 122, 124, and 126 may include any desired area relative to an environment of vehicle 200, for example. In some embodiments, one or more of image capture devices 122, 124, and 126 may be configured to acquire image data from an environment in front of vehicle 200, behind vehicle 200, to the sides of vehicle 200, or combinations thereof.

[0092] Further, the focal length associated with each image capture device 122,124, and/or 126 may be selectable (e.g., by inclusion of appropriate lenses etc.) such that each device acquires images of objects at a desired distance range relative to vehicle 200. For example, in some embodiments image capture devices 122, 124, and 126 may acquire images of close-up objects within a few meters from the vehicle. Image capture devices 122, 124, and 126 may also be configured to acquire images of objects at ranges more distant from the vehicle (e.g., 25 m, 50 m, 100 m, 150 m, or more). Further, the focal lengths of image capture devices 122, 124, and 126 may be selected such that one image capture device (e.g., image capture device 122) can acquire images of objects relatively close to the vehicle (e.g., within 10 m or within 20 m) while the other image capture devices (e.g., image capture devices 124 and 126) can acquire images of more distant objects (e.g., greater than 20 m, 50 m, 100 m, 150 m, etc.) from vehicle 200.

[0093] According to some embodiments, the FOV of one or more image capture devices 122, 124, and 126 may have a wide angle. For example, it may be advantageous to have a FOV of 140 degrees, especially for image capture devices 122, 124, and 126 that may be used to capture images of the area in the vicinity of vehicle 200. For example, image capture device 122 may be used to capture images of the area to the right or left of vehicle 200 and, in such embodiments, it may be desirable for image capture device 122 to have a wide FOV (e.g., at least 140 degrees).

[0094] The field of view associated with each of image capture devices 122, 124, and 126 may depend on the respective focal lengths. For example, as the focal length increases, the corresponding field of view decreases.

[0095] Image capture devices 122, 124, and 126 may be configured to have any suitable fields of view. In one particular example, image capture device 122 may have a horizontal FOV of 46 degrees, image capture device 124 may have a horizontal FOV of

Tesla EX1002 Page 1609 IPR2025-00943

23 degrees, and image capture device 126 may have a horizontal FOV in between 23 and 46 degrees. In another instance, image capture device 122 may have a horizontal FOV of 52 degrees, image capture device 124 may have a horizontal FOV of 26 degrees, and image capture device 126 may have a horizontal FOV in between 26 and 52 degrees. In some embodiments, a ratio of the FOV of image capture device 122 to the FOVs of image capture device 124 and/or image capture device 126 may vary from 1.5 to 2.0. In other embodiments, this ratio may vary between 1.25 and 2.25.

[0096] System 100 may be configured so that a field of view of image capture device 122 overlaps, at least partially or fully, with a field of view of image capture device 124 and/or image capture device 126. In some embodiments, system 100 may be configured such that the fields of view of image capture devices 124 and 126, for example, fall within (e.g., are narrower than) and share a common center with the field of view of image capture device 122. In other embodiments, the image capture devices 122, 124, and 126 may capture adjacent FOVs or may have partial overlap in their FOVs. In some embodiments, the fields of view of image capture devices 122, 124, and 126 may be aligned such that a center of the narrower FOV image capture devices 124 and/or 126 may be located in a lower half of the field of view of the wider FOV device 122.

[0097] FIG. 2F is a diagrammatic representation of exemplary vehicle control systems, consistent with the disclosed embodiments. As indicated in FIG. 2F , vehicle 200 may include throttling system 220, braking system 230, and steering system 240. System 100 may provide inputs (e.g., control signals) to one or more of throttling system 220, braking system 230, and steering system 240 over one or more data links (e.g., any wired and/or wireless link or links for transmitting data). For example, based on analysis of images acquired by image capture devices 122, 124, and/or 126, system 100 may provide control signals to one or more of throttling system 220, braking system 230, and steering system 240 to navigate vehicle 200 (e.g., by causing an acceleration, a turn, a lane shift, etc.). Further, system 100 may receive inputs from one or more of throttling system 220, braking system 230, and steering system 24 indicating operating conditions of vehicle 200 (e.g., speed, whether vehicle 200 is braking and/or turning, etc.). Further details are provided in connection with FIGS. 4-7 , below.

[0098] As shown in FIG. 3A , vehicle 200 may also include a user interface 170 for interacting with a driver or a passenger of vehicle 200. For example, user interface 170 in a vehicle application may include a touch screen 320, knobs 330, buttons 340, and a microphone 350. A driver or passenger of vehicle 200 may also use handles (e.g., located on or near the steering column of vehicle 200 including, for example, turn signal handles), buttons (e.g., located on the steering wheel of vehicle 200), and the like, to interact with system 100. In some embodiments, microphone 350 may be positioned adjacent to a rearview mirror 310. Similarly, in some embodiments, image capture device 122 may be located near rearview mirror 310. In some embodiments, user interface 170 may also include one or more speakers 360 (e.g., speakers of a vehicle audio system). For example, system 100 may provide various notifications (e.g., alerts) via speakers 360.

[0099] FIGS. 3B-3D are illustrations of an exemplary camera mount 370 configured to be positioned behind a rearview mirror (e.g., rearview mirror 310) and against a vehicle windshield, consistent with disclosed embodiments. As shown in FIG. 3B , camera mount 370 may include image capture devices 122, 124, and 126. Image capture devices 124 and 126 may be positioned behind a glare shield 380, which may be flush

Tesla EX1002 Page 1610 IPR2025-00943

against the vehicle windshield and include a composition of film and/or anti-reflective materials. For example, glare shield 380 may be positioned such that it aligns against a vehicle windshield having a matching slope. In some embodiments, each of image capture devices 122, 124, and 126 may be positioned behind glare shield 380, as depicted, for example, in FIG. 3D . The disclosed embodiments are not limited to any particular configuration of image capture devices 122, 124, and 126, camera mount 370, and glare shield 380. FIG. 3C is an illustration of camera mount 370 shown in FIG. 3B from a front perspective.

[0100] As will be appreciated by a person skilled in the art having the benefit of this disclosure, numerous variations and/or modifications may be made to the foregoing disclosed embodiments. For example, not all components are essential for the operation of system 100. Further, any component may be located in any appropriate part of system 100 and the components may be rearranged into a variety of configurations while providing the functionality of the disclosed embodiments. Therefore, the foregoing configurations are examples and, regardless of the configurations discussed above, system 100 can provide a wide range of functionality to analyze the surroundings of vehicle 200 and navigate vehicle 200 in response to the analysis.

[0101] As discussed below in further detail and consistent with various disclosed embodiments, system 100 may provide a variety of features related to autonomous driving and/or driver assist technology. For example, system 100 may analyze image data, position data (e.g., GPS location information), map data, speed data, and/or data from sensors included in vehicle 200. System 100 may collect the data for analysis from, for example, image acquisition unit 120, position sensor 130, and other sensors. Further, system 100 may analyze the collected data to determine whether or not vehicle 200 should take a certain action, and then automatically take the determined action without human intervention. For example, when vehicle 200 navigates without human intervention, system 100 may automatically control the braking, acceleration, and/or steering of vehicle 200 (e.g., by sending control signals to one or more of throttling system 220, braking system 230, and steering system 240). Further, system 100 may analyze the collected data and issue warnings and/or alerts to vehicle occupants based on the analysis of the collected data.

[0102] Further, consistent with disclosed embodiments, the functionality provided by system 100 may cause vehicle 200 to take different actions to navigate vehicle 200 within a lane and/or relative to other vehicles and/or objects. For example, system 100 may adjust the positioning of vehicle 200 relative to a lane within which vehicle 200 is traveling and/or relative to objects positioned near vehicle 200, select a particular lane for vehicle 200 to use while traveling, and take action in response to an encroaching vehicle, such as a vehicle attempting to move into the lane within which vehicle 200 is traveling. Additionally, system 100 may control the speed of vehicle 200 in different scenarios, such as when vehicle 200 is making a turn. System 100 may further cause vehicle 200 to mimic the actions of a leading vehicle or monitor a target vehicle and navigate vehicle 200 so that it passes the target vehicle. Additional details regarding the various embodiments that are provided by system 100 are provided below.

Forward-Facing Multi-Imaging System

[0103] As discussed above, system 100 may provide drive assist functionality that uses a multi-camera system. The multi-camera system may use one or more cameras facing

in the forward direction of a vehicle. In other embodiments, the multi-camera system may include one or more cameras facing to the side of a vehicle or to the rear of the vehicle. In one embodiment, for example, system 100 may use a two-camera imaging system, where a first camera and a second camera (e.g., image capture devices 122 and 124) may be positioned at the front and/or the sides of a vehicle (e.g., vehicle 200). The first camera may have a field of view that is greater than, less than, or partially overlapping with, the field of view of the second camera. In addition, the first camera may be connected to a first image processor to perform monocular image analysis of images provided by the first camera, and the second camera may be connected to a second image processor to perform monocular image analysis of images provided by the second camera. The outputs (e.g., processed information) of the first and second image processors may be combined. In some embodiments, the second image processor may receive images from both the first camera and second camera to perform stereo analysis. In another embodiment, system 100 may use a three-camera imaging system where each of the cameras has a different field of view. Such a system may, therefore, make decisions based on information derived from objects located at varying distances both forward and to the sides of the vehicle. References to monocular image analysis may refer to instances where image analysis is performed based on images captured from a single point of view (e.g., from a single camera). Stereo image analysis may refer to instances where image analysis is performed based on two or more images captured with one or more variations of an image capture parameter. For example, captured images suitable for performing stereo image analysis may include images captured: from two or more different positions, from different fields of view, using different focal lengths, along with parallax information, etc.

[0104] For example, in one embodiment, system 100 may implement a three camera configuration using image capture devices 122-126. In such a configuration, image capture device 122 may provide a narrow field of view (e.g., 34 degrees, or other values selected from a range of about 20 to 45 degrees, etc.), image capture device 124 may provide a wide field of view (e.g., 150 degrees or other values selected from a range of about 100 to about 180 degrees), and image capture device 126 may provide an intermediate field of view (e.g., 46 degrees or other values selected from a range of about 35 to about 60 degrees). In some embodiments, image capture device 126 may act as a main or primary camera. Image capture devices 122-126 may be positioned behind rearview mirror 310 and positioned substantially side-by-side (e.g., 6 cm apart). Further, in some embodiments, as discussed above, one or more of image capture devices 122-126 may be mounted behind glare shield 380 that is flush with the windshield of vehicle 200. Such shielding may act to minimize the impact of any reflections from inside the car on image capture devices 122-126.

[0105] In another embodiment, as discussed above in connection with FIGS. 3B and 3C , the wide field of view camera (e.g., image capture device 124 in the above example) may be mounted lower than the narrow and main field of view cameras (e.g., image devices 122 and 126 in the above example). This configuration may provide a free line of sight from the wide field of view camera. To reduce reflections, the cameras may be mounted close to the windshield of vehicle 200, and may include polarizers on the cameras to damp reflected light.

[0106] A three camera system may provide certain performance characteristics. For example, some embodiments may include an ability to validate the detection of objects by one camera based on detection results from another camera. In the three camera configuration discussed above, processing unit 110 may include, for example, three

processing devices (e.g., three EyeQ series of processor chips, as discussed above), with each processing device dedicated to processing images captured by one or more of image capture devices 122-126.

[0107] In a three camera system, a first processing device may receive images from both the main camera and the narrow field of view camera, and perform vision processing of the narrow FOV camera to, for example, detect other vehicles, pedestrians, lane marks, traffic signs, traffic lights, and other road objects. Further, the first processing device may calculate a disparity of pixels between the images from the main camera and the narrow camera and create a 3D reconstruction of the environment of vehicle 200. The first processing device may then combine the 3D reconstruction with 3D map data or with 3D information calculated based on information from another camera.

[0108] The second processing device may receive images from main camera and perform vision processing to detect other vehicles, pedestrians, lane marks, traffic signs, traffic lights, and other road objects. Additionally, the second processing device may calculate a camera displacement and, based on the displacement, calculate a disparity of pixels between successive images and create a 3D reconstruction of the scene (e.g., a structure from motion). The second processing device may send the structure from motion based 3D reconstruction to the first processing device to be combined with the stereo 3D images.

[0109] The third processing device may receive images from the wide FOV camera and process the images to detect vehicles, pedestrians, lane marks, traffic signs, traffic lights, and other road objects. The third processing device may further execute additional processing instructions to analyze images to identify objects moving in the image, such as vehicles changing lanes, pedestrians, etc.

[0110] In some embodiments, having streams of image-based information captured and processed independently may provide an opportunity for providing redundancy in the system. Such redundancy may include, for example, using a first image capture device and the images processed from that device to validate and/or supplement information obtained by capturing and processing image information from at least a second image capture device.

[0111] In some embodiments, system 100 may use two image capture devices (e.g., image capture devices 122 and 124) in providing navigation assistance for vehicle 200 and use a third image capture device (e.g., image capture device 126) to provide redundancy and validate the analysis of data received from the other two image capture devices. For example, in such a configuration, image capture devices 122 and 124 may provide images for stereo analysis by system 100 for navigating vehicle 200, while image capture device 126 may provide images for monocular analysis by system 100 to provide redundancy and validation of information obtained based on images captured from image capture device 122 and/or image capture device 124. That is, image capture device 126 (and a corresponding processing device) may be considered to provide a redundant sub-system for providing a check on the analysis derived from image capture devices 122 and 124 (e.g., to provide an automatic emergency braking (AEB) system).

[0112] One of skill in the art will recognize that the above camera configurations, camera placements, number of cameras, camera locations, etc., are examples only.

Tesla EX1002 Page 1613 IPR2025-00943

These components and others described relative to the overall system may be assembled and used in a variety of different configurations without departing from the scope of the disclosed embodiments. Further details regarding usage of a multi-camera system to provide driver assist and/or autonomous vehicle functionality follow below.

[0113] FIG. 4 is an exemplary functional block diagram of memory 140 and/or 150, which may be stored/programmed with instructions for performing one or more operations consistent with the disclosed embodiments. Although the following refers to memory 140, one of skill in the art will recognize that instructions may be stored in memory 140 and/or 150.

[0114] As shown in FIG. 4 , memory 140 may store a monocular image analysis module 402, a stereo image analysis module 404, a velocity and acceleration module 406, and a navigational response module 408. The disclosed embodiments are not limited to any particular configuration of memory 140. Further, application processor 180 and/or image processor 190 may execute the instructions stored in any of modules 402-408 included in memory 140. One of skill in the art will understand that references in the following discussions to processing unit 110 may refer to application processor 180 and image processor 190 individually or collectively. Accordingly, steps of any of the following processes may be performed by one or more processing devices.

[0115] In one embodiment, monocular image analysis module 402 may store instructions (such as computer vision software) which, when executed by processing unit 110, performs monocular image analysis of a set of images acquired by one of image capture devices 122, 124, and 126. In some embodiments, processing unit 110 may combine information from a set of images with additional sensory information (e.g., information from radar) to perform the monocular image analysis. As described in connection with FIGS. 5A-5D below, monocular image analysis module 402 may include instructions for detecting a set of features within the set of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, hazardous objects, and any other feature associated with an environment of a vehicle. Based on the analysis, system 100 (e.g., via processing unit 110) may cause one or more navigational responses in vehicle 200, such as a turn, a lane shift, a change in acceleration, and the like, as discussed below in connection with navigational response module 408.

[0116] In one embodiment, stereo image analysis module 404 may store instructions (such as computer vision software) which, when executed by processing unit 110, performs stereo image analysis of first and second sets of images acquired by a combination of image capture devices selected from any of image capture devices 122, 124, and 126. In some embodiments, processing unit 110 may combine information from the first and second sets of images with additional sensory information (e.g., information from radar) to perform the stereo image analysis. For example, stereo image analysis module 404 may include instructions for performing stereo image analysis based on a first set of images acquired by image capture device 124 and a second set of images acquired by image capture device 126. As described in connection with FIG. 6 below, stereo image analysis module 404 may include instructions for detecting a set of features within the first and second sets of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, hazardous objects, and the like. Based on the analysis, processing unit 110 may cause one or more navigational responses in vehicle 200, such as a turn, a lane shift, a change in acceleration, and the like, as discussed below in connection with

Tesla EX1002 Page 1614 IPR2025-00943

navigational response module 408.

[0117] In one embodiment, velocity and acceleration module 406 may store software configured to analyze data received from one or more computing and electromechanical devices in vehicle 200 that are configured to cause a change in velocity and/or acceleration of vehicle 200. For example, processing unit 110 may execute instructions associated with velocity and acceleration module 406 to calculate a target speed for vehicle 200 based on data derived from execution of monocular image analysis module 402 and/or stereo image analysis module 404. Such data may include, for example, a target position, velocity, and/or acceleration, the position and/or speed of vehicle 200 relative to a nearby vehicle, pedestrian, or road object, position information for vehicle 200 relative to lane markings of the road, and the like. In addition, processing unit 110 may calculate a target speed for vehicle 200 based on sensory input (e.g., information from radar) and input from other systems of vehicle 200, such as throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. Based on the calculated target speed, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to trigger a change in velocity and/or acceleration by, for example, physically depressing the brake or easing up off the accelerator of vehicle 200.

[0118] In one embodiment, navigational response module 408 may store software executable by processing unit 110 to determine a desired navigational response based on data derived from execution of monocular image analysis module 402 and/or stereo image analysis module 404. Such data may include position and speed information associated with nearby vehicles, pedestrians, and road objects, target position information for vehicle 200, and the like. Additionally, in some embodiments, the navigational response may be based (partially or fully) on map data, a predetermined position of vehicle 200, and/or a relative velocity or a relative acceleration between vehicle 200 and one or more objects detected from execution of monocular image analysis module 402 and/or stereo image analysis module 404. Navigational response module 408 may also determine a desired navigational response based on sensory input (e.g., information from radar) and inputs from other systems of vehicle 200, such as throttling system 220, braking system 230, and steering system 240 of vehicle 200. Based on the desired navigational response, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and steering system 240 of vehicle 200 to trigger a desired navigational response by, for example, turning the steering wheel of vehicle 200 to achieve a rotation of a predetermined angle. In some embodiments, processing unit 110 may use the output of navigational response module 408 (e.g., the desired navigational response) as an input to execution of velocity and acceleration module 406 for calculating a change in speed of vehicle 200.

[0119] FIG. 5A is a flowchart showing an exemplary process 500A for causing one or more navigational responses based on monocular image analysis, consistent with disclosed embodiments. At step 510, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120. For instance, a camera included in image acquisition unit 120 (such as image capture device 122 having field of view 202) may capture a plurality of images of an area forward of vehicle 200 (or to the sides or rear of a vehicle, for example) and transmit them over a data connection (e.g., digital, wired, USB, wireless, Bluetooth, etc.) to processing unit 110. Processing unit 110 may execute monocular image analysis module 402 to analyze the plurality of images at step 520, as described in further detail in connection with FIGS. 5B-5D below. By performing the analysis,

Tesla EX1002 Page 1615 IPR2025-00943

processing unit 110 may detect a set of features within the set of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, and the like.

[0120] Processing unit 110 may also execute monocular image analysis module 402 to detect various road hazards at step 520, such as, for example, parts of a truck tire, fallen road signs, loose cargo, small animals, and the like. Road hazards may vary in structure, shape, size, and color, which may make detection of such hazards more challenging. In some embodiments, processing unit 110 may execute monocular image analysis module 402 to perform multi-frame analysis on the plurality of images to detect road hazards. For example, processing unit 110 may estimate camera motion between consecutive image frames and calculate the disparities in pixels between the frames to construct a 3D-map of the road. Processing unit 110 may then use the 3D-map to detect the road surface, as well as hazards existing above the road surface.

[0121] At step 530, processing unit 110 may execute navigational response module 408 to cause one or more navigational responses in vehicle 200 based on the analysis performed at step 520 and the techniques as described above in connection with FIG. 4 . Navigational responses may include, for example, a turn, a lane shift, a change in acceleration, and the like. In some embodiments, processing unit 110 may use data derived from execution of velocity and acceleration module 406 to cause the one or more navigational responses. Additionally, multiple navigational responses may occur simultaneously, in sequence, or any combination thereof. For instance, processing unit 110 may cause vehicle 200 to shift one lane over and then accelerate by, for example, sequentially transmitting control signals to steering system 240 and throttling system 220 of vehicle 200. Alternatively, processing unit 110 may cause vehicle 200 to brake while at the same time shifting lanes by, for example, simultaneously transmitting control signals to braking system 230 and steering system 240 of vehicle 200.

[0122] FIG. 5B is a flowchart showing an exemplary process 500B for detecting one or more vehicles and/or pedestrians in a set of images, consistent with disclosed embodiments. Processing unit 110 may execute monocular image analysis module 402 to implement process 500B. At step 540, processing unit 110 may determine a set of candidate objects representing possible vehicles and/or pedestrians. For example, processing unit 110 may scan one or more images, compare the images to one or more predetermined patterns, and identify within each image possible locations that may contain objects of interest (e.g., vehicles, pedestrians, or portions thereof). The predetermined patterns may be designed in such a way to achieve a high rate of "false hits" and a low rate of "misses." For example, processing unit 110 may use a low threshold of similarity to predetermined patterns for identifying candidate objects as possible vehicles or pedestrians. Doing so may allow processing unit 110 to reduce the probability of missing (e.g., not identifying) a candidate object representing a vehicle or pedestrian.

[0123] At step 542, processing unit 110 may filter the set of candidate objects to exclude certain candidates (e.g., irrelevant or less relevant objects) based on classification criteria. Such criteria may be derived from various properties associated with object types stored in a database (e.g., a database stored in memory 140). Properties may include object shape, dimensions, texture, position (e.g., relative to vehicle 200), and the like. Thus, processing unit 110 may use one or more sets of criteria to reject false candidates from the set of candidate objects.

Tesla EX1002 Page 1616 IPR2025-00943

IPR2025-00943
Tesla EX1002 Page 1617

[0124] At step 544, processing unit 110 may analyze multiple frames of images to determine whether objects in the set of candidate objects represent vehicles and/or pedestrians. For example, processing unit 110 may track a detected candidate object across consecutive frames and accumulate frame-by-frame data associated with the detected object (e.g., size, position relative to vehicle 200, etc.). Additionally, processing unit 110 may estimate parameters for the detected object and compare the object's frame-by-frame position data to a predicted position.

[0125] At step 546, processing unit 110 may construct a set of measurements for the detected objects. Such measurements may include, for example, position, velocity, and acceleration values (relative to vehicle 200) associated with the detected objects. In some embodiments, processing unit 110 may construct the measurements based on estimation techniques using a series of time-based observations such as Kalman filters or linear quadratic estimation (LQE), and/or based on available modeling data for different object types (e.g., cars, trucks, pedestrians, bicycles, road signs, etc.). The Kalman filters may be based on a measurement of an object's scale, where the scale measurement is proportional to a time to collision (e.g., the amount of time for vehicle 200 to reach the object). Thus, by performing steps 540-546, processing unit 110 may identify vehicles and pedestrians appearing within the set of captured images and derive information (e.g., position, speed, size) associated with the vehicles and pedestrians. Based on the identification and the derived information, processing unit 110 may cause one or more navigational responses in vehicle 200, as described in connection with FIG. 5A , above.

[0126] At step 548, processing unit 110 may perform an optical flow analysis of one or more images to reduce the probabilities of detecting a "false hit" and missing a candidate object that represents a vehicle or pedestrian. The optical flow analysis may refer to, for example, analyzing motion patterns relative to vehicle 200 in the one or more images associated with other vehicles and pedestrians, and that are distinct from road surface motion. Processing unit 110 may calculate the motion of candidate objects by observing the different positions of the objects across multiple image frames, which are captured at different times. Processing unit 110 may use the position and time values as inputs into mathematical models for calculating the motion of the candidate objects. Thus, optical flow analysis may provide another method of detecting vehicles and pedestrians that are nearby vehicle 200. Processing unit 110 may perform optical flow analysis in combination with steps 540-546 to provide redundancy for detecting vehicles and pedestrians and increase the reliability of system 100.

[0127] FIG. 5C is a flowchart showing an exemplary process 500C for detecting road marks and/or lane geometry information in a set of images, consistent with disclosed embodiments. Processing unit 110 may execute monocular image analysis module 402 to implement process 500C. At step 550, processing unit 110 may detect a set of objects by scanning one or more images. To detect segments of lane markings, lane geometry information, and other pertinent road marks, processing unit 110 may filter the set of objects to exclude those determined to be irrelevant (e.g., minor potholes, small rocks, etc.). At step 552, processing unit 110 may group together the segments detected in step 550 belonging to the same road mark or lane mark. Based on the grouping, processing unit 110 may develop a model to represent the detected segments, such as a mathematical model.

[0128] At step 554, processing unit 110 may construct a set of measurements associated with the detected segments. In some embodiments, processing unit 110

may create a projection of the detected segments from the image plane onto the real-world plane. The projection may be characterized using a 3rd-degree polynomial having coefficients corresponding to physical properties such as the position, slope, curvature, and curvature derivative of the detected road. In generating the projection, processing unit 110 may take into account changes in the road surface, as well as pitch and roll rates associated with vehicle 200. In addition, processing unit 110 may model the road elevation by analyzing position and motion cues present on the road surface. Further, processing unit 110 may estimate the pitch and roll rates associated with vehicle 200 by tracking a set of feature points in the one or more images.

[0129] At step 556, processing unit 110 may perform multi-frame analysis by, for example, tracking the detected segments across consecutive image frames and accumulating frame-by-frame data associated with detected segments. As processing unit 110 performs multi-frame analysis, the set of measurements constructed at step 554 may become more reliable and associated with an increasingly higher confidence level. Thus, by performing steps 550-556, processing unit 110 may identify road marks appearing within the set of captured images and derive lane geometry information. Based on the identification and the derived information, processing unit 110 may cause one or more navigational responses in vehicle 200, as described in connection with FIG. 5A , above.

[0130] At step 558, processing unit 110 may consider additional sources of information to further develop a safety model for vehicle 200 in the context of its surroundings. Processing unit 110 may use the safety model to define a context in which system 100 may execute autonomous control of vehicle 200 in a safe manner. To develop the safety model, in some embodiments, processing unit 110 may consider the position and motion of other vehicles, the detected road edges and barriers, and/or general road shape descriptions extracted from map data (such as data from map database 160). By considering additional sources of information, processing unit 110 may provide redundancy for detecting road marks and lane geometry and increase the reliability of system 100.

[0131] FIG. 5D is a flowchart showing an exemplary process 500D for detecting traffic lights in a set of images, consistent with disclosed embodiments. Processing unit 110 may execute monocular image analysis module 402 to implement process 500D. At step 560, processing unit 110 may scan the set of images and identify objects appearing at locations in the images likely to contain traffic lights. For example, processing unit 110 may filter the identified objects to construct a set of candidate objects, excluding those objects unlikely to correspond to traffic lights. The filtering may be done based on various properties associated with traffic lights, such as shape, dimensions, texture, position (e.g., relative to vehicle 200), and the like. Such properties may be based on multiple examples of traffic lights and traffic control signals and stored in a database. In some embodiments, processing unit 110 may perform multi-frame analysis on the set of candidate objects reflecting possible traffic lights. For example, processing unit 110 may track the candidate objects across consecutive image frames, estimate the real-world position of the candidate objects, and filter out those objects that are moving (which are unlikely to be traffic lights). In some embodiments, processing unit 110 may perform color analysis on the candidate objects and identify the relative position of the detected colors appearing inside possible traffic lights.

[0132] At step 562, processing unit 110 may analyze the geometry of a junction. The

IPR2025-00943
Tesla EX1002 Page 1618

IPR2025-00943
Tesla EX1002 Page 1619

analysis may be based on any combination of: (i) the number of lanes detected on either side of vehicle 200, (ii) markings (such as arrow marks) detected on the road, and (iii) descriptions of the junction extracted from map data (such as data from map database 160). Processing unit 110 may conduct the analysis using information derived from execution of monocular analysis module 402. In addition, Processing unit 110 may determine a correspondence between the traffic lights detected at step 560 and the lanes appearing near vehicle 200.

[0133] As vehicle 200 approaches the junction, at step 564, processing unit 110 may update the confidence level associated with the analyzed junction geometry and the detected traffic lights. For instance, the number of traffic lights estimated to appear at the junction as compared with the number actually appearing at the junction may impact the confidence level. Thus, based on the confidence level, processing unit 110 may delegate control to the driver of vehicle 200 in order to improve safety conditions. By performing steps 560-564, processing unit 110 may identify traffic lights appearing within the set of captured images and analyze junction geometry information. Based on the identification and the analysis, processing unit 110 may cause one or more navigational responses in vehicle 200, as described in connection with FIG. 5A , above.

[0134] FIG. 5E is a flowchart showing an exemplary process 500E for causing one or more navigational responses in vehicle 200 based on a vehicle path, consistent with the disclosed embodiments. At step 570, processing unit 110 may construct an initial vehicle path associated with vehicle 200. The vehicle path may be represented using a set of points expressed in coordinates (x, z),and the distancedibetween two points in the set of points may fall in the range of 1 to 5 meters. In one embodiment, processing unit 110 may construct the initial vehicle path using two polynomials, such as left and right road polynomials. Processing unit 110 may calculate the geometric midpoint between the two polynomials and offset each point included in the resultant vehicle path by a predetermined offset (e.g., a smart lane offset), if any (an offset of zero may correspond to travel in the middle of a lane). The offset may be in a direction perpendicular to a segment between any two points in the vehicle path. In another embodiment, processing unit 110 may use one polynomial and an estimated lane width to offset each point of the vehicle path by half the estimated lane width plus a predetermined offset (e.g., a smart lane offset).

[0135] At step 572, processing unit 110 may update the vehicle path constructed at step 570. Processing unit 110 may reconstruct the vehicle path constructed at step 570 using a higher resolution, such that the distancedkbetween two points in the set of points representing the vehicle path is less than the distancedidescribed above. For example, the distancedkmay fall in the range of 0.1 to 0.3 meters. Processing unit 110 may reconstruct the vehicle path using a parabolic spline algorithm, which may yield a cumulative distance vectorScorresponding to the total length of the vehicle path (i.e., based on the set of points representing the vehicle path).

[0136] At step 574, processing unit 110 may determine a look-ahead point (expressed in coordinates as (xi , zi)) based on the updated vehicle path constructed at step 572. Processing unit 110 may extract the look-ahead point from the cumulative distance vectorS, and the look-ahead point may be associated with a look-ahead distance and look-ahead time. The look-ahead distance, which may have a lower bound ranging from 10 to 20 meters, may be calculated as the product of the speed of vehicle 200 and the look-ahead time. For example, as the speed of vehicle 200 decreases, the look-ahead distance may also decrease (e.g., until it reaches the lower bound). The look-

ahead time, which may range from 0.5 to 1.5 seconds, may be inversely proportional to the gain of one or more control loops associated with causing a navigational response in vehicle 200, such as the heading error tracking control loop. For example, the gain of the heading error tracking control loop may depend on the bandwidth of a yaw rate loop, a steering actuator loop, car lateral dynamics, and the like. Thus, the higher the gain of the heading error tracking control loop, the lower the look-ahead time.

[0137] At step 576, processing unit 110 may determine a heading error and yaw rate command based on the look-ahead point determined at step 574. Processing unit 110 may determine the heading error by calculating the arctangent of the look-ahead point, e.g., arctan (xi/zi). Processing unit 110 may determine the yaw rate command as the product of the heading error and a high-level control gain. The high-level control gain may be equal to: (2 / look-ahead time), if the look-ahead distance is not at the lower bound. Otherwise, the high-level control gain may be equal to: (2 * speed of vehicle 200 / look-ahead distance).

[0138] FIG. 5F is a flowchart showing an exemplary process 500F for determining whether a leading vehicle is changing lanes, consistent with the disclosed embodiments. At step 580, processing unit 110 may determine navigation information associated with a leading vehicle (e.g., a vehicle traveling ahead of vehicle 200). For example, processing unit 110 may determine the position, velocity (e.g., direction and speed), and/or acceleration of the leading vehicle, using the techniques described in connection with FIGS. 5A and 5B , above. Processing unit 110 may also determine one or more road polynomials, a look-ahead point (associated with vehicle 200), and/or a snail trail (e.g., a set of points describing a path taken by the leading vehicle), using the techniques described in connection with FIG. 5E , above.

[0139] At step 582, processing unit 110 may analyze the navigation information determined at step 580. In one embodiment, processing unit 110 may calculate the distance between a snail trail and a road polynomial (e.g., along the trail). If the variance of this distance along the trail exceeds a predetermined threshold (for example, 0.1 to 0.2 meters on a straight road, 0.3 to 0.4 meters on a moderately curvy road, and 0.5 to 0.6 meters on a road with sharp curves), processing unit 110 may determine that the leading vehicle is likely changing lanes. In the case where multiple vehicles are detected traveling ahead of vehicle 200, processing unit 110 may compare the snail trails associated with each vehicle. Based on the comparison, processing unit 110 may determine that a vehicle whose snail trail does not match with the snail trails of the other vehicles is likely changing lanes. Processing unit 110 may additionally compare the curvature of the snail trail (associated with the leading vehicle) with the expected curvature of the road segment in which the leading vehicle is traveling. The expected curvature may be extracted from map data (e.g., data from map database 160), from road polynomials, from other vehicles' snail trails, from prior knowledge about the road, and the like. If the difference in curvature of the snail trail and the expected curvature of the road segment exceeds a predetermined threshold, processing unit 110 may determine that the leading vehicle is likely changing lanes.

[0140] In another embodiment, processing unit 110 may compare the leading vehicle's instantaneous position with the look-ahead point (associated with vehicle 200) over a specific period of time (e.g., 0.5 to 1.5 seconds). If the distance between the leading vehicle's instantaneous position and the look-ahead point varies during the specific period of time, and the cumulative sum of variation exceeds a predetermined threshold (for example, 0.3 to 0.4 meters on a straight road, 0.7 to 0.8 meters on a moderately

IPR2025-00943
Tesla EX1002 Page 1620

curvy road, and 1.3 to 1.7 meters on a road with sharp curves), processing unit 110 may determine that the leading vehicle is likely changing lanes. In another embodiment, processing unit 110 may analyze the geometry of the snail trail by comparing the lateral distance traveled along the trail with the expected curvature of the snail trail. The expected radius of curvature may be determined according to the calculation: (δx<2>+ δx<2>) / 2 / (δx), where δxrepresents the lateral distance traveled and δzrepresents the longitudinal distance traveled. If the difference between the lateral distance traveled and the expected curvature exceeds a predetermined threshold (e.g., 500 to 700 meters), processing unit 110 may determine that the leading vehicle is likely changing lanes. In another embodiment, processing unit 110 may analyze the position of the leading vehicle. If the position of the leading vehicle obscures a road polynomial (e.g., the leading vehicle is overlaid on top of the road polynomial), then processing unit 110 may determine that the leading vehicle is likely changing lanes. In the case where the position of the leading vehicle is such that, another vehicle is detected ahead of the leading vehicle and the snail trails of the two vehicles are not parallel, processing unit 110 may determine that the (closer) leading vehicle is likely changing lanes.

[0141] At step 584, processing unit 110 may determine whether or not leading vehicle 200 is changing lanes based on the analysis performed at step 582. For example, processing unit 110 may make the determination based on a weighted average of the individual analyses performed at step 582. Under such a scheme, for example, a decision by processing unit 110 that the leading vehicle is likely changing lanes based on a particular type of analysis may be assigned a value of "1" (and "0" to represent a determination that the leading vehicle is not likely changing lanes). Different analyses performed at step 582 may be assigned different weights, and the disclosed embodiments are not limited to any particular combination of analyses and weights.

[0142] FIG. 6 is a flowchart showing an exemplary process 600 for causing one or more navigational responses based on stereo image analysis, consistent with disclosed embodiments. At step 610, processing unit 110 may receive a first and second plurality of images via data interface 128. For example, cameras included in image acquisition unit 120 (such as image capture devices 122 and 124 having fields of view 202 and 204) may capture a first and second plurality of images of an area forward of vehicle 200 and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive the first and second plurality of images via two or more data interfaces. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0143] At step 620, processing unit 110 may execute stereo image analysis module 404 to perform stereo image analysis of the first and second plurality of images to create a 3D map of the road in front of the vehicle and detect features within the images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, road hazards, and the like. Stereo image analysis may be performed in a manner similar to the steps described in connection with FIGS. 5A-5D , above. For example, processing unit 110 may execute stereo image analysis module 404 to detect candidate objects (e.g., vehicles, pedestrians, road marks, traffic lights, road hazards, etc.) within the first and second plurality of images, filter out a subset of the candidate objects based on various criteria, and perform multi-frame analysis, construct measurements, and determine a confidence level for the remaining candidate objects. In performing the steps above, processing unit 110 may consider information from both the first and second plurality of images, rather than information from one set of images

Tesla EX1002 Page 1621 IPR2025-00943

alone. For example, processing unit 110 may analyze the differences in pixel-level data (or other data subsets from among the two streams of captured images) for a candidate object appearing in both the first and second plurality of images. As another example, processing unit 110 may estimate a position and/or velocity of a candidate object (e.g., relative to vehicle 200) by observing that the object appears in one of the plurality of images but not the other or relative to other differences that may exist relative to objects appearing if the two image streams. For example, position, velocity, and/or acceleration relative to vehicle 200 may be determined based on trajectories, positions, movement characteristics, etc. of features associated with an object appearing in one or both of the image streams.

[0144] At step 630, processing unit 110 may execute navigational response module 408 to cause one or more navigational responses in vehicle 200 based on the analysis performed at step 620 and the techniques as described above in connection with FIG. 4 . Navigational responses may include, for example, a turn, a lane shift, a change in acceleration, a change in velocity, braking, and the like. In some embodiments, processing unit 110 may use data derived from execution of velocity and acceleration module 406 to cause the one or more navigational responses. Additionally, multiple navigational responses may occur simultaneously, in sequence, or any combination thereof.

[0145] FIG. 7 is a flowchart showing an exemplary process 700 for causing one or more navigational responses based on an analysis of three sets of images, consistent with disclosed embodiments. At step 710, processing unit 110 may receive a first, second, and third plurality of images via data interface 128. For instance, cameras included in image acquisition unit 120 (such as image capture devices 122, 124, and 126 having fields of view 202, 204, and 206) may capture a first, second, and third plurality of images of an area forward and/or to the side of vehicle 200 and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive the first, second, and third plurality of images via three or more data interfaces. For example, each of image capture devices 122, 124, 126 may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0146] At step 720, processing unit 110 may analyze the first, second, and third plurality of images to detect features within the images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, road hazards, and the like. The analysis may be performed in a manner similar to the steps described in connection with FIGS. 5A-5D and 6 , above. For instance, processing unit 110 may perform monocular image analysis (e.g., via execution of monocular image analysis module 402 and based on the steps described in connection with FIGS. 5A-5D , above) on each of the first, second, and third plurality of images. Alternatively, processing unit 110 may perform stereo image analysis (e.g., via execution of stereo image analysis module 404 and based on the steps described in connection with FIG. 6 , above) on the first and second plurality of images, the second and third plurality of images, and/or the first and third plurality of images. The processed information corresponding to the analysis of the first, second, and/or third plurality of images may be combined. In some embodiments, processing unit 110 may perform a combination of monocular and stereo image analyses. For example, processing unit 110 may perform monocular image analysis (e.g., via execution of monocular image analysis module 402) on the first plurality of images and stereo image analysis (e.g., via execution of stereo image

Tesla EX1002 Page 1622 IPR2025-00943

analysis module 404) on the second and third plurality of images. The configuration of image capture devices 122, 124, and 126-including their respective locations and fields of view 202, 204, and 206-may influence the types of analyses conducted on the first, second, and third plurality of images. The disclosed embodiments are not limited to a particular configuration of image capture devices 122, 124, and 126, or the types of analyses conducted on the first, second, and third plurality of images.

[0147] In some embodiments, processing unit 110 may perform testing on system 100 based on the images acquired and analyzed at steps 710 and 720. Such testing may provide an indicator of the overall performance of system 100 for certain configurations of image capture devices 122, 124, and 126. For example, processing unit 110 may determine the proportion of "false hits" (e.g., cases where system 100 incorrectly determined the presence of a vehicle or pedestrian) and "misses."

[0148] At step 730, processing unit 110 may cause one or more navigational responses in vehicle 200 based on information derived from two of the first, second, and third plurality of images. Selection of two of the first, second, and third plurality of images may depend on various factors, such as, for example, the number, types, and sizes of objects detected in each of the plurality of images. Processing unit 110 may also make the selection based on image quality and resolution, the effective field of view reflected in the images, the number of captured frames, the extent to which one or more objects of interest actually appear in the frames (e.g., the percentage of frames in which an object appears, the proportion of the object that appears in each such frame, etc.), and the like.

[0149] In some embodiments, processing unit 110 may select information derived from two of the first, second, and third plurality of images by determining the extent to which information derived from one image source is consistent with information derived from other image sources. For example, processing unit 110 may combine the processed information derived from each of image capture devices 122, 124, and 126 (whether by monocular analysis, stereo analysis, or any combination of the two) and determine visual indicators (e.g., lane markings, a detected vehicle and its location and/or path, a detected traffic light, etc.) that are consistent across the images captured from each of image capture devices 122, 124, and 126. Processing unit 110 may also exclude information that is inconsistent across the captured images (e.g., a vehicle changing lanes, a lane model indicating a vehicle that is too close to vehicle 200, etc.). Thus, processing unit 110 may select information derived from two of the first, second, and third plurality of images based on the determinations of consistent and inconsistent information.

[0150] Navigational responses may include, for example, a turn, a lane shift, a change in acceleration, and the like. Processing unit 110 may cause the one or more navigational responses based on the analysis performed at step 720 and the techniques as described above in connection with FIG. 4 . Processing unit 110 may also use data derived from execution of velocity and acceleration module 406 to cause the one or more navigational responses. In some embodiments, processing unit 110 may cause the one or more navigational responses based on a relative position, relative velocity, and/or relative acceleration between vehicle 200 and an object detected within any of the first, second, and third plurality of images. Multiple navigational responses may occur simultaneously, in sequence, or any combination thereof.

IPR2025-00943
Tesla EX1002 Page 1623

Smart Lane Offset

[0151] System 100 may provide driver assist functionality that monitors and adjusts the free space between vehicle 200 and boundaries (e.g., lane boundaries) within which vehicle 200 is traveling. A lane may refer to a designated or intended travel path of a vehicle and may have marked (e.g., lines on a road) or unmarked boundaries (e.g., an edge of a road, a road barrier, guard rail, parked vehicles, etc.). For instance, as a default, system 100 may maximize or increase the free space between vehicle 200 and the current lane boundaries. Further, under certain conditions, system 100 may offset vehicle 200 toward one side of the current lane of travel without leaving the lane. For example, where vehicle 200 is traveling around a curve, system 100 may decrease the offset towards the inside of the curve (e.g., adjust the position of vehicle 200 so that it is closer to the inside of the curve). As another example, where obstacles-such as parked cars, pedestrians, or cyclists-are present on one side of the lane, system 100 may increase the offset on the side of vehicle 200 where the obstacles are located (e.g., adjust the position of vehicle 200 so that it is further from the obstacles).

[0152] FIG. 8A illustrates vehicle 200 traveling on a roadway 800 in which the disclosed systems and methods for identifying lane constraints and operating vehicle 200 within the lane constraints may be used. As shown, vehicle 200 may have a first vehicle side 802, which may be a first distance 805 from a first lane constraint 810. Similarly, vehicle 200 may have a second vehicle side 812 opposite from first vehicle side 802, and second vehicle side 812 may be a second distance 815 from second lane constraint 820. In this manner, first lane constraint 810 and second lane constraint 820 may define a lane within which vehicle 200 may travel.

[0153] Processing unit 110 may be configured to determine first lane constraint 810 and second lane constraint 820 based on a plurality of images acquired by image capture device 122-126 that processing unit 110 may receive via data interface 128. According to some embodiments, first lane constraint 810 and/or second lane constraint 820 may be identified by visible lane boundaries, such as dashed or solid lines marked on a road surface. Additionally or alternatively, first lane constraint 810 and/or second lane constraint 820 may include an edge of a road surface or a barrier. Additionally or alternatively, first lane constraint 810 and/or second lane constrain 820 may include markers (e.g., Botts' dots). According to some embodiments, processing unit 110 may determine first lane constraint 810 and/or second lane constraint 820 by identifying a midpoint of a road surface width. According to some embodiments, if processing unit 110 identifies only first lane constraint 810, processing unit 110 may estimate second lane constraint 820, such as based on an estimated lane width or road width. Processing unit 110 may identify lane constraints in this manner when, for example, lines designating road lanes are not painted or otherwise labeled.

[0154] Detection of first lane constraint 810 and/or second lane constraint 820 may include processing unit 110 determining their 3D models in camera coordinate system. For example, the 3D models of first lane constraint 810 and/or second lane constraint 820 may be described by a third-degree polynomial. In addition to 3D modeling of first lane constraint 810 and/or second lane constraint 820, processing unit 110 may perform multi-frame estimation of host motion parameters, such as the speed, yaw and pitch rates, and acceleration of vehicle 200. Optionally, processing unit 110 may detect static and moving vehicles and their position, heading, speed, and acceleration, all relative to vehicle 200. Processing unit 110 may further determine a road elevation model to transform all of the information acquired from the plurality of images into 3D

IPR2025-00943
Tesla EX1002 Page 1624

space.

[0155] Generally, as a default condition, vehicle 200 may travel relatively centered within first and second lane constraints 810 and/or 820. However, in some circumstances, environmental factors may make this undesirable or unfeasible, such as when objects or structures are present on one side of the road. Thus, there may be circumstances in which it may be desirable or advantageous for vehicle 200 to be closer to one lane constraint or the other. Processing unit 110 may also be configured to determine whether such circumstances, called lane offset conditions, exist based on the plurality of images.

[0156] To determine whether lane offset conditions exist, processing unit 110 may be configured to determine the presence of one or more objects 825 in the vicinity of vehicle 200. For example object 825 may comprise another vehicle, such as a car, truck or motorcycle traveling on roadway 800. Processing unit 110 may determine, for each object 825, an offset profile. The offset profile may include a determination of whether in its current and predicted position object 825 will be within a predefined range of vehicle 200. If processing unit 110 determines that object 825 is or will be within a predefined range of vehicle 200, processing unit 110 may determine whether there is enough space within first lane constraint 810 and second lane constraint 820, for vehicle 200 to bypass object 825. If there is not enough space, processing unit 110 may execute a lane change. If there is enough space, processing unit 110 may determine whether there is enough time to bypass object 825 before the distance between vehicle 200 and object 825 closes. If there is enough time, processing unit 110 may initiate and or activate the offset maneuver to bypass object 825. The offset maneuver may be executed by processing unit 110 so that the movement of vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 to 0.75 m/sec. The offset maneuver may be executed so that the maximum amplitude of the offset maneuver will occur when the gap between vehicle 200 and object 825 closes. If there is not enough time, processing unit 110 may initiate braking (e.g., by transmitting electronic signals to braking system 230) and then initiate offset maneuver.

[0157] After each offset profile is set for each object 825, processing unit may combine offset profiles such that the maximum lateral distance is kept from each object 825 that vehicle 200 is bypassing. That is, offset maneuvers may be modified and/or executed such that the lateral distance between vehicle 200 and each object 825 that vehicle 200 is bypassing is maximized.

[0158] Additionally or alternatively, object 825 may be a parked or stationary vehicle, a wall, or a person, such as pedestrian or cycling traffic. Processing unit 110 may determine a lane offset condition exists based on certain characteristics of object 825. For example, processing unit 110 may determine object 825 constitutes a lane offset condition if the height of object 825 exceeds a predetermined threshold, such as 10 cm from a road surface. In this manner, processing unit 110 may be configured to filter out objects that would not present a lane offset condition, such as small pieces of debris.

[0159] Additionally or alternatively, processing unit 110 may be configured to determine whether a lane offset condition exists on first vehicle side 802 based on a position of a target object 830 on vehicle first side 802. Target object 830 may include one or more objects that constitute a lane offset condition. Additionally or alternatively, target object 830 may not constitute a lane offset condition. For example, target object 830 may be a

Tesla EX1002 Page 1625 IPR2025-00943

street sign or other object a sufficient distance from the travel lane or occurring with sufficient infrequency to not constitute a lane offset condition. Processing unit 110 may be configured to conclude a lane offset condition does not exist if such a target object 830 is detected.

[0160] Processing unit 110 may be configured to cause vehicle 200 to travel such that first distance 805 and second distance 810 are substantially different to avoid or stay farther away from object 825. For example, if a lane offset condition exists on first vehicle side 802, processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is greater than second distance 815. This is illustrated, for example, in FIG. 8A . Alternatively, if a lane offset condition exists on second vehicle side 812, processing unit 110 may create or execute an offset profile that is configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is less than second distance 815. This is illustrated, for example, in FIG 8B .

[0161] As another example, a parked vehicle may be located on one side of vehicle 200 and a pedestrian may be located on the other side of vehicle 200. In such a situation, system 100 may determine that is desirable to reduce the distance between vehicle 200 and the parked vehicle, in order to provide more space between vehicle 200 and the pedestrian. Accordingly, processing unit 110 may cause vehicle 200 to travel such that the distance between the parked vehicle is reduced.

[0162] The identity of first vehicle side 802 and second vehicle side 812 need not be static, such that at different points in time, processing unit 110 may identify first vehicle side 802 as the driver's side or as the passenger's side of vehicle 200. For example, the identity of first vehicle side 802 may change depending on the presence of one or more objects 825 on that side of vehicle 200, such that if one lane offset condition exists based on object 825, it exists on first vehicle side 802. In such circumstances, processing unit 110 may create or execute an offset profile that is configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805, which is greater than second distance 815, is provided toward object 825.

[0163] According to some embodiments, when a lane offset condition exists on first vehicle side 802, processing unit 110 may create an offset profile configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is at least 1.25 times greater than second distance 815. Additionally or alternatively, when a lane offset condition exists on first vehicle side 802, processing unit 110 may create an offset profile that is configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is at least 2 times greater than second distance 815.

[0164] In some embodiments, as shown in FIG. 8C , processing unit 110 may determine a lane offset condition exists on both first vehicle side 802 and on second vehicle side 812. Processing unit 110 may then create an offset profile for this offset condition that is configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is substantially the same as second distance 815.

[0165] Processing unit 110 may be configured so that the offset profiles for moving

Tesla EX1002 Page 1626 IPR2025-00943

objects 825 as well as stationary objects 825 are prioritized. For example, processing unit 110 may prioritize offset profiles to keep a predefined distance between vehicle 200 and moving objects 825 and/or static objects 825. Processing unit 200 may prioritize these offset conditions over those offset conditions based on certain lane characteristics, such as a curve. The last priority lane offset condition may include mimicking a lead vehicle.

[0166] Processing unit 110 may factor offset profiles or offset conditions into a desired look-ahead point, which may be located at the middle of the lane defined by first lane constraint 810 and second lane constraint 820. Processing unit 110 may compute and implement a yaw rate command based on the look-ahead point and the offset profiles. The yaw rate command may be used to compute a torque command based on the yaw rate command and the yaw rate of vehicle 200 measured by using a gyroscope. Processing unit 110 may transmit the torque command to a steering system (e.g., steering system 240) of vehicle 200.

[0167] According to some embodiments, input from a driver of vehicle 200, such as via user interface 170, may override processing unit 110. For example, a user input may override the control of processing unit 110 to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is greater than second distance 815.

[0168] Additionally or alternatively, processing unit 110 may be configured to determine, based on the plurality of images captured by image capture device 122, whether a lane offset condition is substantially non-existent. This may include circumstances in which one or more objects 825 below a certain size threshold are detected, one or more objects 825 detected are outside of a predetermined distance from the lane constraint, and/or when no objects 825 are detected. Under such circumstances, processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is equal to or substantially equal to second distance 815.

[0169] A lane offset condition may include a curved road, as illustrated in FIG. 8D . Processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is toward the outside of the curved road relative to vehicle 200 and second distance 815 is toward an inside of the curved road relative to vehicle 200. For example, first distance 805 may be less than second distance 815.

[0170] FIG. 9 is an exemplary block diagram of memory 140 and/or 150, which may stare instructions for performing one or more operations consistent with disclosed embodiments. As illustrated in FIG. 9 , memory 140 may store one or more modules for performing the lane offset condition detection and responses described herein. For example, memory 140 may store a lane constraint module 910 and a lane offset module 920. Application processor 180 and/or image processor 190 may execute the instructions stored in any of modules 910 and 920 included in memory 140. One of skill in the art will understand that references in the following discussions to processing unit 110 may refer to application processor 180 and image processor 190 individually or collectively. Accordingly, steps of any of the following processes may be performed by one or more processing devices.

[0171] Lane constraint module 910 may store instructions which, when executed by

Tesla EX1002 Page 1627 IPR2025-00943

processing unit 110, may detect and define first lane constraint 810 and second lane constraint 820. For example, lane offset module 910 may process the plurality of images received from at least one image capture device 122-124 to detect first lane constraint 810 and second lane constraint 820. As discussed above, this may include identifying painted lane lines and/or measuring a midpoint of a road surface.

[0172] Lane offset module 920 may store instructions which, when executed by processing unit 110, may detect the presence of a lane offset condition and/or identify whether lane offset condition is present on first vehicle side 802 and/or second vehicle side 812. For example, lane offset module 920 may process the plurality of images to detect the presence of object(s) 825 on first vehicle side 802 and/or second vehicle side 812. Lane offset module 920 may process the plurality of images to detect a curve in the road on which vehicle 200 is traveling. Additionally or alternatively, lane offset module 920 may receive information from another module or other system indicative of the presence of object(s) 825 and/or a curved lane. Lane offset module 920 may execute control to change the position of vehicle 200 to first lane constraint 810 and/or second lane constraint 820. For example, if lane offset module 920 determines a lane offset condition only on first vehicle side 802, lane offset module 920 may move vehicle 200 closer to second lane constraint 820. On the other hand, if lane offset module 920 determines a lane offset condition only on first vehicle side 812, lane offset module 920 may move vehicle 200 closer to first lane constraint 810.

[0173] Fig. 10 illustrates a process 1000 for navigating vehicle 200, consistent with disclosed embodiments. Process 1000 may identify lane constraints that define a lane of travel for vehicle travel. Process 1000 include causing vehicle 200 to travel closer to one or the other lane constraint, such as first lane constraint 810 or second lane constraint 820, in response to detecting a lane offset condition. In this manner, process 1000 may be used to drive vehicle 200 within lane constraints as well as account for environmental factors that may make it more desirable to shift vehicle 200 within first lane constraint 810 and second lane constraint 820.

[0174] At step 1010, process 1000 may include acquiring, using at least one image capture device 122, 124, and/or 126, a plurality of images of an area in the vicinity of vehicle 200. For example, processing unit 110 may receive the plurality of images may through data interface 128. For example, processing unit 110 may be configured to determine the presence of one or more objects 825 on the side of the road. For example, object 825 may be a parked or stationary vehicle, a wall, or a person, such as pedestrian or cycling traffic. Processing unit 110 may determine if a lane offset condition exists based on certain characteristics of object 825. For example, processing unit 110 may determine that object 825 constitutes a lane offset condition if the height of object 825 exceeds a predetermined threshold, such as 10 cm from a road surface. In this manner, processing unit 110 may be configured to filter out objects that would not present a lane offset condition, such as small pieces of debris.

[0175] Additionally or alternatively, processing unit 110 may be configured to determine whether a lane offset condition exists on first vehicle side 802 based on a position of a target object 830 on vehicle first side 802. Target object 830 may be one or more objects that constitute a lane offset condition. Additionally or alternatively, target object 830 may not constitute a lane offset condition. For example, target object 830 may be a road barrier or a street sign. Processing unit 110 may be configured to conclude a lane offset condition does not exist if such a target object 830 is detected.

IPR2025-00943
Tesla EX1002 Page 1628

[0176] At step 1020, process 1000 may determine from the plurality of images a first lane constraint on first vehicle side 802 and a second lane constraint on second vehicle side 812. For example, processing unit 110 may be configured to determine first lane constraint 810 and second lane constraint 820 based on the plurality of images received via data interface 128.

[0177] At step 1030, process 1000 may determine whether a lane offset condition exists on first vehicle side 802. If a lane offset condition does exist on first vehicle side 802, at step 1040, process 1000 may include causing vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that second distance 815 is less than first distance 805. Processing unit 110 may be configured to cause vehicle 200 to travel such that first distance 805 and second distance 810 are substantially different to avoid or stay farther away from object 825. For example, if a lane offset condition exists on first vehicle side 802, processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is greater than second distance 815. This is illustrated, for example, in FIG. 8A . Alternatively, if a lane offset condition exists on second vehicle side 812, processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is less than second distance 815. This is illustrated, for example, in FIG 8B . For example, to navigate vehicle 200 according to process 1000, processing unit 110 may transmit electronic signals to one or more of throttling system 220, braking system 230, and/or steering system 240 of vehicle 200.

[0178] If a lane offset condition does not exist on first vehicle side 802, in step 1050, method may include determining whether lane offset condition exists on second vehicle side 812. If a lane offset condition exists on second vehicle side 812, the method may include causing vehicle 200 to travel between first lane constraint 810 and second lane constraint 820 such that first distance 805 is less than second distance 815.

Navigating a Vehicle to a Default Lane

[0179] System 100 may provide driver assist functionality that monitors the location of vehicle 200 in a current lane of travel and moves vehicle 200 into a predetermined default travel lane if the current lane of travel is not the predetermined default travel lane. The predetermined default travel lane may be defined by a user input and/or as a result of processing images received from at least one image capture device 122-126.

[0180] System 100 may also provide notification of an impending lane change of vehicle 200 to the predetermined default travel lane to the driver and passengers of vehicle 200 and/or other drivers. For example, system 100 may activate a turn signal positioned at or near one of bumper regions 210. Additionally or alternatively, system 100 may sound an audible notification, such as through speakers 360. Vehicle 200 may use system 100 to select a particular lane as predetermined default travel lane when vehicle is traveling. When the current lane within which vehicle 200 is traveling-e.g., the current lane of travel-differs from the predetermined default travel lane, system 100 may cause vehicle 200 to make a lane change, e.g., into the predetermined default travel lane.

[0181] FIG. 11 illustrates vehicle 200 traveling in a current lane of travel 1110 of roadway 800. According to some embodiments, vehicle 200 may have a predetermined default travel lane 1120. For example, predetermined default travel lane 1120 may be

IPR2025-00943
Tesla EX1002 Page 1629

the right-most lane of roadway 800. Navigation system 100 may be configured to determine whether vehicle 200 is traveling in predetermined default travel lane 1120 and, if not, cause vehicle 200 to change to predetermined default travel lane 1120. For example, processing unit 110 may be configured to compare current lane of travel 1110 with predetermined default travel lane 1120 and cause vehicle 200 to return to predetermined default travel lane 1120.

[0182] According to some embodiments, processing unit 110 may detect an object 1130 in predetermined default travel lane 1120 based on a plurality of images acquired by one or more of image capture devices 122-126. Processing unit 110 may detect the position and speed of vehicle 200 relative to object 1130 by analyzing the images using techniques described in connection with FIGS. 5A-5D , above. According to some embodiments, processing unit 110 may determine whether vehicle 200 should pass object 1130. This may be desirable if, for example, the speed and/or acceleration of vehicle 200 relative to object 1130 exceeds a certain threshold.

[0183] FIG. 12A is an exemplary block diagram of memory 140, which may store instructions for performing one or more operations consistent with disclosed embodiments. As illustrated in FIG. 12A , memory 140 may store one or more modules for performing the default travel lane and current travel lane detection and responses described herein. For example, memory 140 may store a current lane detection module 1210, a default lane detection module 1220, and an action module 1230.

[0184] Current lane detection module 1210 may store instructions which, when executed by processing unit 110, may detect and define current lane of travel 1110. For example, processing unit 110 may execute current lane detection module 1212 to process the plurality of images received from at least one image capture device 122-124 and detect current lane of travel 1110. This may include detecting the relative position of vehicle 200 with respect to the edge of roadway 800, a midpoint of roadway 800, and/or lane markers, such as painted lane lines. Processing unit 110 may detect the position of vehicle 200 relative to roadway 800 by analyzing the images using techniques described in connection with FIGS. 5A-5D , above.

[0185] Default lane detection module 1220 may store instructions which, when executed by processing unit 110, may identify predetermined default travel lane 1120 and compare default lane detection module 1220 to current lane of travel 1110 identified by current lane detection module 1210. Default lane detection module 1220 may identify predetermined default travel lane 1120 based upon a user input and/or the plurality of images received from at least one image capture device 122-124. For example, a user may select a default lane through voice commands received by speakers 360 or by making a selection on a menu displayed on touch screen 320.

[0186] In some embodiments, predetermined default travel lane 1120 may be the right-most travel lane, or it may be the lane that has the more desirable features. For example, default lane detection module 1220 may include instructions to select the lane with less traffic, the lane in which vehicles are traveling closest to speed of vehicle 200, and/or a lane closest to an approaching exit or street to which vehicle 200 will turn. In some embodiments, predetermined default travel lane 1120 may be determined based on data received from position sensor 130, such as a global positioning system.

[0187] Action module 1230 may store instructions which, when executed by processing unit 110, may cause a response to the identification of current lane of travel 1110 and

IPR2025-00943
Tesla EX1002 Page 1630

predetermined default travel lane 1120. For example, processing unit 110 may execute action module 1230 to cause vehicle 200 to change lanes if current lane of travel 1110 is not the same as predetermined default travel lane 1120. Additionally or alternatively, action module 1230 may initiate a notification of the lane change, such as activating turn signal 210 and or causing an audible announcement.

[0188] Action module 1230 may include instructions that cause processing unit 110 to determine whether vehicle 200 should move out of predetermined default travel lane 1120. For example, action module 1230 may determine that object 1130 in predetermined default travel lane 1120 should be bypassed by vehicle 200.

[0189] Action module 1230 may further include instructions that cause processing unit 110 to determine whether it is safe to make a lane change. For example, prior to executing a lane change, processing unit 110 may evaluate determine one or more characteristics of roadway 800 based on, for example, analysis of the plurality of images conducted using techniques described in connection with FIGS. 5A-5D , above. If the analysis indicates conditions are not safe to change lanes, processing unit 110 may determine that vehicle 200 should stay in the current lane.

[0190] Action module 1230 may include additional instructions for manipulating vehicle 200 to maneuver out of and return to predetermined default travel lane 1120, such as for bypassing object 1130. When executing these instructions, processing unit 110 may estimate a location of predetermined default travel lane 1120 using estimated motion of vehicle 200. This may be advantageous in circumstances in which predetermined default travel lane 1120 does not remain detectable by image capture devices 122-124.

[0191] Processing unit 110 may be configured to execute an offset profile that offsets vehicle 200 from predetermined default travel lane 1120 by a full lane width, which may be estimated based on the distance between detected lane markers. The offset maneuver defined by offset profile may be executed by processing unit 110 so that the movement of vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 to 0.75 m/sec. During the offset maneuver, processing unit 110 may execute instructions from action module 1230 that switch to following the center of current lane of travel 1110. These instructions may include gradually transitioning from following the predicted lane with offset to following current lane of travel 1110. A gradual transition may be accomplished by defining desired look-ahead point so that it is not allowed to change in time more than the recommended slew rate, which may be around 0.25 m/sec.

[0192] The instructions for transitioning vehicle 200 from predetermined default travel lane 1120 may be used to transition vehicle 200 back into predetermined default travel lane. For example, processing unit 110 may be configured to execute an offset profile that offsets vehicle 200 from current lane of travel 1110 by a full lane width, which may be estimated based on the distance between detected lane markers. The offset maneuver defined by offset profile may be executed by processing unit 110 so that the movement of vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 to 0.75 m/sec. During the offset maneuver, processing unit 110 may execute instructions from action module 1230 that switch to following the center of predetermined default travel lane 1120. These instructions may include gradually transitioning from following the predicted lane with offset to following predetermined default travel lane 1120. A gradual transition may be accomplished by defining desired look-ahead point so that it is not allowed to change in time more than

IPR2025-00943
Tesla EX1002 Page 1631

the recommended slew rate, which may be around 0.25 m/sec.

[0193] At one or more of the steps for transitioning vehicle 200 between two lanes, processing unit 110 may compute and implement a yaw rate command based on the look-ahead point and the offset profiles. The yaw rate command may be used to compute a torque command based on the yaw rate command and the yaw rate of vehicle 200 measured using a gyroscope. Processing unit 110 may transmit a torque command to a steering system (e.g., steering system 240) of vehicle 200.

[0194] FIG. 12B illustrates a process 1250 for navigating vehicle 200 to a default lane consistent with disclosed embodiments. According to some embodiments, process 1250 may be implemented by one or more components of navigation system 100, such as at least one processing unit 110. Process 1250 may identify current lane of travel 1110 of vehicle 200 and compare it with predetermined default travel lane 1120 to determine whether vehicle 200 is traveling within predetermined default travel lane 1120. If vehicle 200 is not traveling in predetermined default travel lane 1120, processing unit 110 may cause vehicle 200 to change lanes so that it navigates to default travel lane 1120.

[0195] At step 1260, one or more of image capture devices 122-124 may acquire a plurality of images of an area in a vicinity of vehicle 200. Processing unit 110 may receive the plurality of images via data interface 128.

[0196] At step 1270, processing unit 110 may execute current lane detection module 1210 detect and define current lane of travel 1110. For example, processing unit 110 may process the plurality of images to, for example, detect the relative position of vehicle 200 with respect to the edge of roadway 100, a midpoint of roadway 100, and/or painted lane markers. Processing unit 110 may detect the position of vehicle 200 relative to roadway 800 by analyzing the images using techniques described in connection with FIGS. 5A-5D , above.

[0197] At step 1280, processing unit 110 may execute default lane detection module 1220 to identify predetermined default travel lane 1120 and compare default lane detection module 1220 to current lane of travel 1110 identified by current lane detection module 1210. Default lane detection module 1220 may identify predetermined default travel lane 1120 based upon a user input and/or the plurality of images received from at least one image capture device 122-124.

[0198] For example, as discussed above, processing unit 110 may determine predetermined default travel lane 1120 based on any number of factors. Predetermined default travel lane 1120 may be the right-most lane among the plurality of travel lanes 1110 and 1120. Additionally or alternatively, predetermined default travel lane 1120 may be determined based on an input received from a user via user interface 170, by processing the plurality of images, and/or based on data received from position sensor 130, such as a global positioning system. Predetermined default travel lane 1120 may be determined once, once per vehicle trip, or on regular intervals. For example, predetermined default travel lane 1120 may be determined dynamically in response to conditions at a given location.

[0199] If current lane 1110 is not the same as predetermined default travel lane 1120, processing unit 110 may execute action module 1230 to cause vehicle 100 to navigate to default travel lane 1120. For example, processing unit 110 may transmit electronic

IPR2025-00943
Tesla EX1002 Page 1632

signals to one or more of throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to trigger the desired response. For example, processing unit 110 may cause steering system 240 to turn the steering wheel of vehicle 200 to achieve a rotation of a predetermined angle.

[0200] Further, prior to executing a lane change, processing unit 110 may determine whether it is safe (e.g., there are no vehicles or objects in the way) to change lanes. As discussed above, processing unit 110 may evaluate one or more characteristics of roadway 800 based on, for example, analysis of the plurality of images conducted using techniques described in connection with FIGS. 5A-5D , above. If conditions are not safe to change lanes, processing unit 110 may determine that vehicle 200 should stay in the current lane.

[0201] In some embodiments, execution of action module 1230 may initiate a notification of the lane change. For example, processing unit 110 may cause a signal to activate a turn signal and/or cause an audible announcement via speakers 360. Accordingly, such notification may occur prior to vehicle 200 changing lanes. In other embodiments, execution of action module 1230 may initiate a notification to a driver of vehicle 200 that vehicle 200 is not traveling in a default lane without causing the lane change to occur. Accordingly, in such an embodiment, the driver of vehicle 200 may determine whether to manually take action and steer vehicle 200 to change lanes.

[0202] In some embodiments, system 100 may notify the user of vehicle 200 and/or other drivers in the vicinity of vehicle 200 that vehicle 200 will be changing lanes if vehicle 200 will be moved to predetermined default travel lane 1120. For example, processing unit 110 may activate a turn signal of vehicle 200 prior to causing vehicle 200 to change lanes. Additionally or alternatively, an audible announcement to notify other drivers in the vicinity may be caused by processing unit 110. The audible announcement may be transmitted via, for example, Bluetooth, to receivers included in other nearby vehicles. Additionally or alternatively, at least one processing unit 110 may cause an audible indicator to notify the driver of vehicle 200 prior to causing vehicle 200 to switch lanes. In this manner, driver of vehicle 200 can anticipate the control system 100 will exert on vehicle 200.

Controlling Velocity of a Turning Vehicle

[0203] System 100 may provide driver assist functionality that controls the velocity (i.e., speed and/or direction) of vehicle 200 in different scenarios, such as while making a turn (e.g., through a curve). For example, system 100 may be configured to use a combination of map, position, and/or visual data to control the velocity of vehicle 200 when turning. In particular, system 100 may control the velocity of vehicle 200 as it approaches a curve, while navigating a curve, and/or as it exits a curve. System 100 may consider different factors to control the velocity depending on the position of vehicle 200 relative to the curve at a given time. For example, system 100 may initially cause vehicle 200 to slow down in response to learning that a curve is located a certain distance ahead of vehicle 200. While vehicle 200 turns through the curve, however, system 100 may cause vehicle 200 to accelerate and/or charge direction based on an analysis of the characteristics of the curve.

[0204] FIGS. 13A and 13B are diagrammatic representations of a vehicle (e.g., vehicle 200) approaching, and navigating, a curve with one or more characteristics on road 1300, consistent with the disclosed embodiments. As illustrated in FIG. 13A , visual

IPR2025-00943 Tesla EX1002 Page 1633

indicators, such as road sign 1302 located near road 1300, may inform vehicle 200 that a curve lies ahead. In response to the warning provided by road sign 1302, a driver of vehicle 200 and/or a system 100 providing driver assist functionality of vehicle 200 may adjust the velocity of vehicle 200 (e.g., by slowing down and/or steering) to safely navigate the curve.

[0205] As vehicle 200 approaches the curve, system 100 may detect the lane width 1306 and the curve radius 1308 by using image capture devices 122 and/or 124 (which may be cameras) of vehicle 200. Based on the curve radius 1308, the lane width 1306, and other characteristics associated with the curve detected using image capture devices 122 and/or 124, system 100 may adjust the velocity of vehicle 200. Specifically, system 100 may adjust the velocity as vehicle 200 approaches the curve and/or while vehicle 200 navigates the curve. System 100 may also adjust the velocity as vehicle 200 exits the curve, in response to, for example, traffic light 1304 appearing after the curve. Further detail regarding adjustments to the velocity of vehicle 200 and detection of the lane width 1306, the radius of curvature or curve radius 1308, and characteristics associated with the curve are described in connection with FIGS. 14 and 15 , below.

[0206] In some embodiments, system 100 may recognize a curve to be navigated based on map data and/or vehicle position information (e.g., GPS data). System 100 may determine an initial target velocity for the vehicle based on one or more characteristics of the curve as reflected in the map data. System 100 may adjust a velocity of the vehicle to the initial target velocity and determine, based on one or more images acquired by one or more of image capture devices 122-126, one or more observed characteristics of the curve. System 100 may determine an updated target velocity based on the one or more observed characteristics of the curve and adjust the velocity of the vehicle to the updated target velocity.

[0207] FIG. 14 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments. As indicated in FIG. 14 , memory 140 may store a curve recognition module 1402, a characteristic observation module 1404, and a velocity module 1406. The disclosed embodiments are not limited to any particular configuration of memory 140. Further, processing unit 110 may execute the instructions stored in any of modules 1402-1406 included in memory 140.

[0208] In one embodiment, curve recognition module 1402 may store software instructions (such as computer vision software) which, when executed by processing unit 110, detects a curve by, for example, analyzing one or more sets of images and/or using map and/or position data (e.g., data stored in map database 160). The image sets may be acquired by one or more of image capture devices 122, 124, and 126. Based on the analysis, system 100 (e.g., via processing unit 110) may cause a change in velocity of vehicle 200. For example, processing unit 110 may cause vehicle 200 to reduce its speed by a predetermined amount and/or adjust its direction by a particular angle if processing unit 110 detects the presence of a curve located within a minimum threshold distance ahead of vehicle 200.

[0209] In one embodiment, characteristic observation module 1404 may store software instructions (such as computer vision software) which, when executed by processing unit 110, observes characteristics associated with a curve detected by system 100 (e.g., via execution of curve recognition module 1402). Processing unit 110 may

Tesla EX1002 Page 1634 IPR2025-00943

observe characteristics associated with a detected curve by analyzing one or more sets of images acquired by one or more of image capture devices 122, 124, and 126 and/or using map and/or position data (e.g., data stored in map database 160). The analysis may yield information associated with the curve, such as curve radius 1308, lane width 1306, a degree of curvature, a rate of change in curvature, a degree of banking, a length or arc length of the curve, and the like. These characteristics may include estimated values (e.g., curve characteristics calculated using preexisting data and/or mathematical models based on such data) and actual values (e.g., curve characteristics calculated by analyzing captured images of the curve). Based on the analysis, processing unit 110 may cause a change in velocity of vehicle 200. For example, processing unit 110 may cause vehicle 200 to reduce its speed and/or adjust its direction by a particular angle in view of a high rate of change in curvature associated with a curve located 25 meters ahead of vehicle 200.

[0210] In one embodiment, velocity module 1406 may store software instructions configured to analyze data from one or more computing and electromechanical devices configured to determine a target velocity and cause a change in velocity of vehicle 200 to the target velocity. For example, processing unit 110 may execute velocity module 1406 to calculate a target velocity for vehicle 200 based on data derived from execution of curve recognition module 1402 and characteristic observation module 1404. Such data may include, for example, an initial target position and initial velocity, an updated target position and updated velocity, a final target position and final velocity, and the like. In addition, processing unit 110 may calculate a velocity for vehicle 200 based on input from other systems of vehicle 200, such as a throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. Based on the calculated target velocity, processing unit 110 may transmit electronic signals (e.g., via a controller area network bus ("CAN bus")) to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to trigger a change in the speed and/or direction of vehicle 200 by, for example, physically depressing the brake, easing up off the accelerator of vehicle 200, or steering vehicle 200 in a particular direction (e.g., at a particular angle). An increase in speed may be associated with a corresponding acceleration (e.g., between 0.2 m/sec<2>and 3.0 m/sec<2>). Conversely, a decrease in speed may be associated with a corresponding deceleration. In addition, the acceleration or deceleration may be based on the road type (e.g., highway, city street, country road, etc.), the presence of any speed constraints nearby (e.g., sharp curves, traffic lights), as well as a road width or lane width 1306. For purposes of this disclosure, deceleration may refer to an acceleration having a negative value.

[0211] FIG. 15 is a flow chart showing an exemplary process 1500 for controlling the velocity of a vehicle based on a detected curve and observed characteristics of the curve and/or map data regarding the curve consistent with disclosed embodiments. At step 1510, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120. For instance, a camera included in image acquisition unit 120 (such as image capture device 122) may capture a plurality of images and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive more than one plurality of images via a plurality of data interfaces. For example, processing unit 110 may receive a plurality of images from each of image capture devices 122, 124, 126, each of which may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

Tesla EX1002 Page 1635 IPR2025-00943

[0212] At step 1520, processing unit 110 may execute curve recognition module 1502 to detect a curve located ahead of vehicle 200 based on, for example, map and/or position data (e.g., data stored in map database 160). Such data may indicate the location of a curve that is present on a given road. The data may indicate the absolute position of a particular curve (e.g., via GPS coordinates) or the relative position of the curve (e.g., by describing the curve relative to vehicle 200, road sign 1302, traffic light 1304, and/or other descriptive landmarks). In some embodiments, processing unit 110 may execute curve recognition module 1502 to detect the curve by analyzing the plurality of images. The analysis may be conducted using techniques described in connection with FIGS. 5A-5D and 6 , above.

[0213] At step 1530, processing unit 110 may execute velocity module 1406 to determine an initial target velocity for vehicle 200 based on, for example, data derived from execution of curve recognition module 1402. The initial target velocity may reflect a speed and a direction for vehicle 200 to safely enter an upcoming curve and may be determined using techniques described in connection with FIG. 14 , above. For example, velocity module 1406 may determine the initial target velocity based on known or estimated characteristics associated with the upcoming curve, such as the curve radius 1308 and the degree of banking. Velocity module 1406 may calculate the initial target velocity in accordance with mathematical models and/or formulae and based on various constraints. In some embodiments, for example, velocity module 1406 may determine the initial target velocity based on the presence of traffic lights (such as traffic light 1304) before, during, or after a curve, weather conditions (e.g., light rain, heavy rain, windy, clear, etc.), road conditions (e.g., paved road, unpaved road, potholes, etc.), posted speed limits, the presence of vehicles or pedestrians nearby, etc. The initial target velocity may also be based on a lateral acceleration limit associated with vehicle 200. In some embodiments, the lateral acceleration limit may be in the range of 0.2 m/sec<2>and 3.0 m/sec<2>and may be adjusted depending on the curve characteristics, the road width, or lane width 1306. Additionally, the initial target velocity may be based on a focal range associated with one or more of cameras, such as image capture devices 122 and 124.

[0214] For example, velocity module 1406 may determine the initial target velocity based on the radius of curvature of the road ahead of vehicle 200. Velocity module 1406 may determine the radius of curvature at various distances from vehicle 200 and store the information in a vectorR(1:n). In one embodiment, velocity module 1406 may determine the radius of curvature at: (1) the current position of vehicle 200, (2) a distanceDendfrom vehicle 200, whereDendrepresents the distance ahead of vehicle 200 at which knowledge of the road characteristics allows system 100 to control navigational responses in vehicle 200 such that the driver may feel comfortable, and (3) at incremental distances from vehicle 200, such as every 1 meter, etc. The value of distanceDendmay depend on the current speed of vehicle 200,Vego .For example, the greater the value ofVego(e.g., the faster vehicle 200 is traveling), the greater the value ofDend(e.g., the longer the distance may be taken into account to provide system 100 with sufficient time to provide a navigational response, such as comfortably reducing the speed of vehicle 200 based on an approaching sharp turn or curve). In one embodiment,Dendmay be calculated according to the following equation:Dend= (Vego<2>/2)/(0.7/40*Vego+0.5). Thus, the vectorD(1:n) = [1:1:n], wheren= floor(Dend).

[0215] In some embodiments, velocity module 1406 may determine the initial target velocity based on a lateral acceleration constraint,amax, e.g., the maximum allowed lateral acceleration,amaxmay fall in the range of 1.5 to 3.5 meters / sec<2>. Based

IPR2025-00943
Tesla EX1002 Page 1636

onamaxand the radius of curvature vectorR(discussed above), velocity module 1406 may limit the speed of vehicle 200 (e.g., at different points or distances from the current position of vehicle 200) to a speedvlim, calculated as:vlim(1:n)=sqrt(amax*R(1:n)). In addition, a deceleration constraint that follows from the determination ofvlimmay be calculated as:declim= min{((vlim(1:n)<2>-Vego<2>)/2), max{dmin,D(1:n)-2*Vego}}, wheredminis equal to 40 meters. The value ofdminmay represent a point at which deceleration feels comfortable for the driver of vehicle 200.

[0216] In some embodiments, velocity module 1406 may determine the initial target velocity based on a curvature slew rate constraint,slewmax, which may be based on the curvature rate of change.slewmaxmay fall in the range of 0.5 to 1.5 meters / sec<3>. Velocity module 1406 may determinevlimaccording the following equation:vlim (1:n) = (slewmax/abs{diff{1/R(1:n)}})<1/3>. Further, velocity module 1406 may determinedeclimaccording to the equation:declim= min{((vlim(1:n)<2>-Vego<2>)/2), max {dmin , D(1:n) - 2 *Vego}}, wheredminis equal to 40 meters.

[0217] In some embodiments, velocity module 1406 may determine a comfort constraint such that the driver of vehicle 200 is comfortable while vehicle 200 is undergoing various navigational responses, such as making a turn or navigating a curve. Velocity module 1406 may determinevlimaccording to the following equation:vlim (1:n) = 4/500*R(1:n) + 22.5. Further, velocity module 1406 may determinedeclimaccording to the equation:declim= min{((vlim(1:n)<2>-Vego<2>)/2), max {dmin , D(1:n) - 2*Vego}}, wheredminis equal to 40 meters.

[0218] Velocity module 1406 may merge the three constraints determined above-the lateral acceleration constraint, the curvature slew rate constraint, and comfort constraint-by using the smallest value ofdeclimcalculated in association with each constraint. After merging the constraints, velocity module 1406 may arrive at the initial target velocity by determining how much to adjust the speed of vehicle 200, denotedvcom, according to the following equation:vcom=vcom(prev) +declim*Δt,wherevcom(prev)represents the prior speed of vehicle 200 (e.g., a previously determined speed, the speed of vehicle 200 prior to any adjustment by velocity module 1406, etc.), and Δtrepresents the time between the determination ofvcomandvcom(prev).

[0219] At step 1540, processing unit 110 may execute velocity module 1406 to adjust a velocity of vehicle 200 to the initial target velocity determined at step 1530. Based on the current velocity of vehicle 200 and the initial target velocity determined at step 1530, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to change the velocity of vehicle 200 to the initial target velocity. Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240. For example, processing unit 110 may transmit electronic signals that cause system 100 to physically depress the brake by a predetermined amount or ease partially off the accelerator of vehicle 200. Further, processing unit 110 may transmit electronic signals that cause system 100 to steer vehicle 200 in a particular direction. In some embodiments, the magnitude of acceleration (or de-acceleration) may be based on the difference between the current velocity of vehicle 200 and the initial target velocity. Where the difference is large (e.g., the current velocity is 10 km/hr greater than the initial target velocity), for example, system 100 may change the current velocity by braking in a manner that maximizes deceleration and results in vehicle 200 achieving the initial target velocity in the shortest amount of time safely possible. Alternatively, system 100

IPR2025-00943
Tesla EX1002 Page 1637

may apply the brakes in a manner that minimizes deceleration and results in vehicle 200 achieving the initial target velocity gradually. More generally, system 100 may change the current velocity to the initial target velocity according to any particular braking profile (e.g., a braking profile calling for a high level of braking for the first two seconds, and a low level of braking for the subsequent three seconds). The disclosed embodiments are not limited to any particular braking profile or manner of braking.

[0220] At step 1550, processing unit 110 may execute characteristic observation module 1404 to determine one or more characteristics of a curve located ahead of vehicle 200 based on, for example, analysis of the plurality of images conducted using techniques described in connection with FIGS. 5A-5D , above. The characteristics may include, for example, estimated and/or actual values associated with curve radius 1308, a road width or lane width 1306, a degree of curvature, a rate of change in curvature, a degree of banking, and the like. The characteristics may also include, for example, information conveyed by one or more road signs (such as road sign 1302). Furthermore, the characteristics may include the length or arc length of the detected curve, which may be a function of the speed of vehicle 200 as it approaches the curve. Processing unit 110 may execute characteristic observation module 1404 to calculate values associated with these characteristics (e.g., curve radius 1308, the degree of curvature, the rate of change in curvature, the arc length of the curve, etc.) based on mathematical models and/or formulae. In some embodiments, processing unit 110 may execute characteristic observation module 1404 to construct a mathematical model expressing the curvature of the curve as a function of distance (e.g., distance traveled). Such a model may be subject to constraints such as a maximum lateral acceleration constraint (e.g., in the range of 1.5 to 3.0 m/sec<2>) and/or a maximum lateral acceleration derivative constraint (e.g., in the range of 0.8 to 1.2 m/sec<2>). These constraints may yield a maximum speed for vehicle 200 as a function of distance traveled.

[0221] At step 1560, processing unit 110 may execute velocity module 1406 to calculate an updated target velocity based on the characteristics determined at step 1550. The updated target velocity may reflect an updated speed and/or direction for vehicle 200 to safely perform any combination of: (i) entering an upcoming curve, (ii) navigating through a curve, or (iii) exiting a curve. The updated target velocity may be determined using techniques described in connection with FIG. 14 and step 1530, above. For example, velocity module 1406 may determine the updated target velocity based on characteristics associated with a curve and mathematical models and/or formulae.

[0222] At step 1570, processing unit 110 may execute velocity module 1406 to adjust the velocity of vehicle 200 to the updated target velocity determined at step 1560. This adjustment may be accomplished using techniques described in connection with step 1540, above. For example, based on the current velocity of vehicle 200 and the updated target velocity determined at step 1560, processing unit 110 may transmit electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to change the velocity of vehicle 200 to the updated target velocity. At this point, the current velocity of vehicle 200 is not necessarily the same as the initial target velocity determined at step 1530. For example, processing unit 110 may have determined the updated target velocity at step 1560 while in the process of adjusting the velocity of vehicle 200 to the initial target velocity. This may be true in the case where the difference in the velocity of vehicle 200 and the initial target velocity is large (e.g., greater than 10 km/hr) and the associated

Tesla EX1002 Page 1638 IPR2025-00943

acceleration is small.

[0223] In some embodiments, system 100 may make regular updates to the target velocity based on continued observations of characteristics of a curve (e.g., based on image data depicting portions of the curve that vehicle 200 is approaching). Additionally, after system 100 has navigated vehicle 200 through a curve, system 100 may cause vehicle 200 to accelerate to a new target velocity suitable for traveling a segment of the road that does not have curves, as determined by analyzing map data, positional data, and/or image data acquired by system 100.

Mimicking A Leading Vehicle

[0224] System 100 may provide driver assist functionality that causes vehicle 200 to mimic (or decline to mimic) a leading vehicle in different scenarios, such as when the leading vehicle switches lanes, accelerates, or makes a turn. For example, system 100 may be configured to detect the leading vehicle by analyzing a plurality of images and determine one or more actions taken by the leading vehicle. In some scenarios, such as when the leading vehicle is turning at an intersection, system 100 may be configured to cause vehicle 200 to decline to mimic the turn. In other scenarios, such as when the leading vehicle is turning at an intersection without changing lanes within the intersection, system 100 may be configured to cause vehicle 200 to mimic the turn. System 100 may also be configured to cause vehicle 200 to mimic one or more actions of the leading vehicle based on a navigation history of the leading vehicle.

[0225] FIGS. 16A and 16B are diagrammatic representations of a primary vehicle 200a mimicking one or more actions of a leading vehicle 200b on road 1600a, consistent with the disclosed embodiments. As illustrated in FIGS. 16A and 16B , primary vehicle 200a may be trailing leading vehicle 200b while traveling in the same lane of road 1600a. The lane may have a left edge 1620, a right edge 1630, and a midpoint 1610, with w1and w2indicating the left and right halves of the lane, respectively. Primary vehicle 200a and leading vehicle 200b may be positioned in the lane such that neither vehicle is centered on the midpoint 1610. As shown in FIG. 16A , for example, primary vehicle 200a may be centered to the right of midpoint 1610 such that the distance c1between primary vehicle 200a and left edge 1620 is greater than the distance c2between primary vehicle 200a and right edge 1630. In contrast, leading vehicle 200b may be centered to the left of midpoint 1610 such that the distance c3between leading vehicle 200b and left edge 1620 is less than the distance c4between leading vehicle 200b and right edge 1630. Different configurations of the positions of primary vehicle 200a and leading vehicle 200b are possible. For example, primary vehicle 200a and leading vehicle 200b may be positioned anywhere on road 1600a. Further, although FIG. 16A depicts leading vehicle 200b as having a camera system, the disclosed embodiments are not limited to a configuration in which leading vehicle 200b includes a camera system.

[0226] System 100 of primary vehicle 200a may cause primary vehicle 200a to mimic one or more actions of leading vehicle 200b. As shown in FIG. 16B , for example, primary vehicle 200a may mimic a lane shift performed by leading vehicle 200b. After mimicking the lane shift, distance c1may be equal to distance c3, and distance c2may be equal to distance c4. In some embodiments, primary vehicle 200a may mimic a lane shift performed by leading vehicle 200b in a way such that distances c1and c3, and c2and c4, are unequal. For example, in the case where primary vehicle 200a mimics a leftward lane shift of primary vehicle 200b and is positioned (after the lane shift) to the

IPR2025-00943
Tesla EX1002 Page 1639

left of primary vehicle 200a, c1will be less than c3and c2will be greater than c4. Alternatively, where primary vehicle 200a mimics a leftward lane shift of primary vehicle 200b and is positioned to the right of primary vehicle 200a, c1will be greater than c3and c2will be less than c4.

[0227] FIGS. 16C and 16D are diagrammatic representations of a primary vehicle 200a mimicking one or more actions of a leading vehicle 200b on road 1600b, consistent with the disclosed embodiments. As illustrated in FIG. 16C , primary vehicle 200a may be traveling behind leading vehicle 200b, which may make a left turn at an intersection. System 100 of primary vehicle 200a may cause primary vehicle 200a to mimic the left turn of leading vehicle 200b, as shown in FIG. 16D . System 100 may cause primary vehicle 200a to mimic the left turn based on various factors, such as whether leading vehicle 200b changes lanes while turning within the intersection, the navigation history associated with leading vehicle 200b, the relative difference in speed between primary vehicle 200a and leading vehicle 200b, and the like. Additionally, primary vehicle 200a may mimic leading vehicle 200b under various conditions, such as based on whether leading vehicle 200b is traveling along the same route or the same portion of a route that primary vehicle 200a is traveling along. Further detail regarding scenarios in which system 100 may cause primary vehicle 200a to mimic one or more actions of leading vehicle 200b are described in connection with FIGS. 17-22 , below.

[0228] In some embodiments, system 100 may be configured to determine whether to mimic a particular movement of a leading vehicle. For example, system 100 may determine that certain movements of the leading vehicle do not need to be mimicked because the movements do not affect the

[0229] course of the leading vehicle. For example, system 100 may determine that small changes in movements, such as moving within a particular lane, do not need to be mimicked. In other embodiments, system 100 may implement a smoothing operation to filter out small movements of a leading vehicle. The smoothing operation may cause the primary vehicle to implement more significant and/or important movements (e.g., changing lanes) of a leading vehicle while filtering out smaller movements not affecting the overall course of the leading vehicle.

[0230] FIG. 17 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments. As shown in FIG. 17 , memory 140 may store a lane constraint module 1702, an action detection module 1704, and an action response module 1706. The disclosed embodiments are not limited to any particular configuration of memory 140. Further, processing unit 110 may execute the instructions stored in any of modules 1702-1706 included in memory 140.

[0231] In one embodiment, lane constraint module 1702 may store software instructions (such as computer vision software) which, when executed by processing unit 110, determine one or more constraints associated with a lane that primary vehicle 200a is traveling in. The lane constraints may include a midpoint 1610, a left edge 1620, and a right edge 1630 of the lane. The lane constraints may also include one or more lines or other symbols marked on the surface of road 1600a or 1600b. Processing unit 110 may execute lane constraint module 1702 to determine the constraints by, for example, analyzing one or more sets of images and/or using map and/or position data (e.g., data stored in map database 160). The image sets may be acquired by one or more of image capture devices 122, 124, and 126. Based on the analysis, system 100

Tesla EX1002 Page 1640 IPR2025-00943

(e.g., via processing unit 110) may cause primary vehicle 200a to travel within a lane defined by left edge 1620 and right edge 1630. In some embodiments, system 100 may cause primary vehicle 200a to move away from left edge 1620 or right edge 1630 if one of distance c1or c2is less than a predetermined threshold.

[0232] In one embodiment, action detection module 1704 may store software instructions (such as computer vision software) which, when executed by processing unit 110, detects one or more actions taken by leading vehicle 200b traveling in front of primary vehicle 200a. Leading vehicle 200b may be traveling in the same or different lane as primary vehicle 200a. Processing unit 110 may execute action detection module 1704 to detect one or more actions taken by leading vehicle 200b by, for example, analyzing one or more sets of images acquired by one or more of image capture devices 122, 124, and 126. Based on the analysis, processing unit 110 may determine, for example, that leading vehicle 200b has shifted lanes, made a turn, accelerated, decelerated, applied its brakes, and/or the like. As another example, leading vehicle 200b may perform a maneuver such that a first offset distance on a side of leading vehicle 200b adjacent to a first lane constraint is different from a second offset distance on a side of leading vehicle 200b adjacent to a second lane constraint. For example, processing unit 110 may perform the analysis based on the techniques described in connection with FIGS. 5A-5D , above.

[0233] In one embodiment, action response module 1706 may store software instructions (such as computer vision software) which, when executed by processing unit 110, determines one or more actions for primary vehicle 200a to take based on one or more actions taken by leading vehicle 200b detected by system 100 (e.g., via execution of action detection module 1704). For example, processing unit 110 may execute action response module 1706 to determine whether to mimic one or more actions of leading vehicle 200b. This determination may be based on the nature of leading vehicle 200b's actions (e.g., turn, lane shift, etc.), information associated with primary vehicle 200a (e.g., speed, distance c1and c2from the lane edges, etc.), road and environmental conditions (e.g., potholes, heavy rain or wind, etc.), and the like. In the case where processing unit 110 determines to mimic one or more actions of leading vehicle 200b, processing unit 110 may accomplish this by, for example, transmitting electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to trigger a turn, lane shift, change in speed, and/or change in direction of primary vehicle 200a. Processing unit 110 may use the techniques described in connection with FIGS. 4-7 , above, to cause primary vehicle 200a to mimic the one or more actions. To mimic the one or more actions, system 100 may provide control signals to one or more of throttling system 220, braking system 230, and steering system 240 to navigate vehicle 200 (e.g., by causing an acceleration, a turn, a lane shift, etc.). Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240.

[0234] FIG. 18 is a flow chart showing an exemplary process for causing primary vehicle 200a to mimic one or more actions of leading vehicle 200b consistent with disclosed embodiments. At step 1810, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120. For instance, a camera included in image acquisition unit 120 (such as image capture device 122) may capture a plurality of images and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive more than one plurality of images via a plurality of data interfaces. For example, processing unit 110 may receive a plurality of

Tesla EX1002 Page 1641 IPR2025-00943

images from each of image capture devices 122, 124, 126, each of which may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0235] At steps 1820 and 1830, processing unit 110 may execute lane constraint module 1702 to determine constraints for each side of primary vehicle 200a associated with a lane that primary vehicle 200a is traveling in. Processing unit 110 may determine the constraints by analyzing the images received at step 1810. Thus, at step 1820, processing unit 110 may determine a first lane constraint on a first side of primary vehicle 200a (e.g., the distance between the left side of primary vehicle 200a and left edge 1620 of the lane, c1) and, at step 1830, a second lane constraint on a second side opposite the first side (e.g., the distance between the right side of primary vehicle 200a and right edge 1620 of the lane, c2). In some embodiments, processing unit 110 may determine the constraints by using map and/or position data (e.g., data stored in map database 160) indicating the position of primary vehicle 200a relative to midpoint 1610, left edge 1620, and right edge 1630 of the lane. Processing unit 110 may determine the constraints based on both an analysis of the images received at step 1810 and map and/or position data; doing so may increase the confidence level associated with the constraints.

[0236] At step 1840, processing unit 110 may execute lane constraint module 1702 to cause primary vehicle 200a to travel within the lane constraints determined at steps 1820 and 1830. For example, the left side of primary vehicle 200a may be positioned to the left of left edge 1620, or the right side of primary vehicle 200a may be positioned to the right of right edge 1630. In such a scenario, processing unit 110 may transmit electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to cause primary vehicle 200a to adjust its position and travel within the lane constraints.

[0237] At step 1850, processing unit 110 may execute action detection module 1704 to locate a leading vehicle 200b within the images received at step 1810 by analyzing the images. Processing unit 110 may detect the presence of vehicles in the images using the techniques described in connection with FIGS. 5A and 5B , above. For each detected vehicle, processing unit 110 may construct a multi-frame model of the detected vehicle's position, velocity (e.g., speed and direction), and acceleration relative to primary vehicle 200a. Processing unit 110 may use the techniques described in connection with FIGS. 5A and 5B , above, to construct such a model. In addition, processing unit 110 may determine a lane of travel associated with each detected vehicle, based on, for example, lane or other road markings and lane constraints determined at steps 1820 and 1830. Processing unit 110 may determine leading vehicle 200b to be the detected vehicle that is closest to, and is traveling in the same lane as, primary vehicle 200a. At step 1860, processing unit 110 may execute action detection module 1704 to determine one or more actions taken by the leading vehicle 200b detected within the images at step 1850. For example, processing unit 110 may determine that leading vehicle 200b has shifted lanes, made a turn, accelerated, decelerated, and/or the like. In addition, processing unit 110 may analyze the images to determine characteristics associated with leading vehicle 200b, such as its speed and position (e.g., relative to primary vehicle 200a). These determinations may be based on the techniques described in connection with FIGS. 5A and 5B , above. In some embodiments, processing unit 110 may analyze the images to determine characteristics associated with an action performed by leading vehicle 200b. For example, where leading vehicle 200b makes a turn, processing unit 110 may determine

position information for leading vehicle 200b right before it begins the turn, position information for leading vehicle 200b as it completes the turn, a turning radius, and a speed profile describing changes in speed (e.g., over time) for leading vehicle 200b as it makes the turn.

[0238] At step 1870, processing unit 110 may execute action response module 1706 to cause primary vehicle 200a to mimic the one or more actions of leading vehicle 200b determined at step 1850. For example, processing unit 110 may transmit electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to trigger a turn, lane shift, change in speed, and/or change in direction of primary vehicle 200a. In addition, processing unit 110 may construct a set of path points to guide changes in position and speed for primary vehicle 200a. The path points may be based on lane marks (e.g., midpoint 1610, left edge 1620, right edge 1630, etc.) represented by a polynomial (e.g., of the third-degree) and target coordinates associated with a target vehicle, such as leading vehicle 200b. Processing unit 110 may adjust the coefficients of the polynomial so that the path points pass through the target coordinates. Processing unit 110 may also offset the path points (e.g., relative to left edge 1620 and right edge 1630) based on a predetermined offset amount. As processing unit 110 adjusts the polynomial coefficients and offsets the path points, processing unit 110 may also calculate local curvature information for each of the resulting segments formed between the path points.

[0239] In some embodiments, processing unit 110 (e.g., via execution of action response module 1706) may cause primary vehicle 200a to mimic leading vehicle 200b within certain position, speed, and/or acceleration constraints associated with the two vehicles. For example, the constraints may require a predetermined minimum distance (e.g., a safety distance) between primary vehicle 200a and leading vehicle 200b or a predetermined maximum difference in speed and/or acceleration between the vehicles. Thus, if mimicking an action performed by leading vehicle 200b would violate one or more constraints (e.g., mimicking an acceleration would result in primary vehicle 200a being too close to leading vehicle 200b), processing unit 110 may cause primary vehicle 200a to decline to mimic the one or more actions of leading vehicle 200b. In addition, even while mimicking an action performed by leading vehicle 200b (e.g., a lane shift), processing unit 110 may cause primary vehicle 200a to follow leading vehicle 200b's speed profile and maintain a predetermined distance or a safety distance from leading vehicle 200b. Processing unit 110 may accomplish this by causing primary vehicle 200a to: (i) continuously close the distance gap from leading vehicle 200b when the gap is greater than the predetermined minimum distance, (ii) continuously close the gap between the speeds of primary vehicle 200a and leading vehicle 200b, and (iii) match the acceleration of leading vehicle 200b to the extent it is consistent with predetermined constraints. To control the extent to which of these actions should be performed, processing unit 110 may associate each action with weights in real-time based on the current and predicted distance to leading vehicle 200b. Thus, for example, as the distance to leading vehicle 200b decreases more weight may be put on decelerating rather than matching the acceleration of leading vehicle 200b.

[0240] As another example, in the case where primary vehicle 200a is approaching leading vehicle 200b at a speed greater than the speed of leading vehicle 200b, processing unit 110 may cause primary vehicle 200a to gradually slow down such that, when the distance between the two vehicles equals a predetermined safety distance

IPR2025-00943
Tesla EX1002 Page 1643

such as a minimum following distance, the speed of the two vehicles is the same. In another example, where a leading vehicle 200b abruptly appears in front of primary vehicle 200a, causing the distance between the two vehicles to be less than the minimum following distance, processing unit 110 may cause primary vehicle 200a to gradually slow down to avoid approaching leading vehicle 200b any closer while at the same increasing the distance between the vehicles until it reaches the minimum following distance.

[0241] FIG. 19 is a flow chart showing an exemplary process 1900 for causing primary vehicle 200a to decline to mimic a turn of leading vehicle 200b, consistent with disclosed embodiments. At step 1910, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 1810 of FIG. 18 , above. At step 1920, processing unit 110 may execute action detection module 1704 to locate a leading vehicle 200b within the images received at step 1910 by analyzing the images, as described in connection with step 1850 of FIG. 18 , above. At step 1925, processing unit 110 may execute action detection module 1704 to determine whether leading vehicle 200b is turning at an intersection, using the techniques described in connection with step 1860 of FIG. 18 , above. If leading vehicle 200b is determined to be turning at the intersection (step 1925, yes), processing unit 110 may cause primary vehicle 200a to decline to mimic the turn of leading vehicle 200b at step 1930. Otherwise (step 1925, no), process 1900 concludes.

[0242] FIG. 20 is a flow chart showing an exemplary process 2000 for causing primary vehicle 200a to mimic or decline a turn of leading vehicle 200b, consistent with disclosed embodiments. At step 2010, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 1810 of FIG. 18 , above. At step 2020, processing unit 110 may execute action detection module 1704 to locate a leading vehicle 200b within the images received at step 2010 by analyzing the images, as described in connection with step 1850 of FIG. 18 , above. At step 2025, processing unit 110 may execute action detection module 1704 to determine whether leading vehicle 200b is turning at an intersection, using the techniques described in connection with step 1860 of FIG. 18 , above. If leading vehicle 200b is determined to be turning at the intersection (step 2025, yes), at step 2035, processing unit 110 may execute action detection module 1704 to determine whether leading vehicle 200b is changing lanes within the intersection, using the techniques described in connection with step 1860 of FIG. 18 , above. Otherwise (step 2025, no), process 2000 concludes. If leading vehicle 200b is determined to be changing lanes within the intersection (step 2035, yes), processing unit 110 may cause primary vehicle 200a to decline to mimic the turn of leading vehicle 200b at step 2050. Otherwise (step 2035, no), processing unit 110 may cause primary vehicle 200a to mimic the turn of leading vehicle 200b at step 2040.

[0243] FIG. 21 is a flow chart showing another exemplary process 2100 for causing primary vehicle 200a to mimic one or more actions of a first leading vehicle, consistent with disclosed embodiments. At step 2110, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 1810 of FIG. 18 , above. At steps 2120 and 2130, processing unit 110 may execute action detection module 1704 to locate a first and second leading vehicle within the images received at step 2110 by analyzing the images, as described in connection with step 1850 of FIG. 18 , above. At step 2135, processing unit 110 may execute action detection module 1704 to determine

IPR2025-00943
Tesla EX1002 Page 1644

whether the first leading vehicle path has a lower turning radius than the second leading vehicle path, using the techniques described in connection with step 1860 of FIG. 18 , above. If the first leading vehicle path is determined to have a lower turning radius than the second leading vehicle path (step 2135, yes), processing unit 110 may cause primary vehicle 200a to mimic the turn of the first leading vehicle at step 2140. Otherwise (step 2135, no), process 2100 concludes.

[0244] FIG. 22 is a flow chart showing an exemplary process 2200 for causing primary vehicle 200a to mimic one or more actions of leading vehicle 200b based on a navigation history, consistent with disclosed embodiments. At step 2210, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 1810 of FIG. 18 , above. At step 2220, processing unit 110 may execute action detection module 1704 to locate a leading vehicle 200b within the images received at step 2210 by analyzing the images, as described in connection with step 1850 of FIG. 18 , above.

[0245] At step 2230, processing unit 110 may determine position information for leading vehicle 200b by analyzing the images received at step 2210. Position information for leading vehicle 200b may indicate the position of leading vehicle 200b relative to primary vehicle 200a. Processing unit 110 may analyze the images to derive position information for leading vehicle 200b and one or more actions taken by leading vehicle 200b by using the techniques described in connection with FIGS. 5A-5D , above.

[0246] At step 2240, processing unit 110 may compare the position information for leading vehicle 200b determined at step 2230 to predetermined map data (e.g., data stored in map database 160). For example, by comparing the position information to predetermined map data, processing unit 110 may derive position coordinates for leading vehicle 200b. In addition, processing unit 110 may use the predetermined map data as an indicator of a confidence level associated with the position information determined at step 2230.

[0247] At step 2250, processing unit 110 may create a navigation history associated with leading vehicle 200b by, for example, tracking the actions taken by leading vehicle 200b. Processing unit 110 may associate the actions taken by leading vehicle 200b (e.g., determined from analysis of the images at step 2230) with predetermined map data from step 2240. Thus, the navigation history may indicate that leading vehicle 200b performed certain actions (e.g., a turn, a lane shift, etc.) at certain position coordinates.

[0248] At step 2260, processing unit 110 may execute action response module 1706 to determine whether to cause primary vehicle 200a to mimic one or more actions of leading vehicle 200b based on the navigation history created at step 2250 and using the techniques described in connection with step 1870 of FIG. 18 , above. For example, processing unit 110 (e.g., via action response module 1706) may cause primary vehicle 200a to mimic a turn of leading vehicle 200b, even where no turn is detected from analysis of the images (e.g., at step 2230), when the navigation history indicates that leading vehicle 200b has successfully navigated at least one prior turn. More generally, processing unit 110 may cause primary vehicle 200a to mimic an action taken by leading vehicle 200b in cases where the navigation history indicates the leading vehicle 200b had successfully performed such an action.

Navigating a Vehicle to Pass Another Vehicle

Tesla EX1002 Page 1645 IPR2025-00943

[0249] System 100 may provide driver assist functionality that monitors the vicinity around vehicle 200 and aborts lane changes in the event that the lane change is deemed to be at risk of causing a collision, such as with another vehicle (e.g., a target vehicle). For example, system 100 may enable vehicle 200 to complete a pass of a target vehicle if the target vehicle is determined to be in a different lane than vehicle 200. If, before completion of a pass, system 100 determines that the target vehicle is entering the lane in which vehicle 200 is traveling, system 100 may cause vehicle 200 to abort the pass. For example, aborting the pass may include causing vehicle 200 to brake or causing vehicle 200 to stop accelerating.

[0250] As illustrated in FIG. 23 , vehicle 200 may travel on roadway 2300. Vehicle 200 may be equipped with system 100 and may implement any of the disclosed embodiments for identifying lane constraints and operating vehicle 200 within the lane constraints. System 100 may also facilitate lane changes made by vehicle 200. For example, system 100 may determine whether circumstances allow vehicle 200 to safely change lanes and, if the circumstances change during the lane change, cause vehicle 200 to abort lane change.

[0251] As shown in FIG. 23 and as previously discussed with respect to FIG. 8 , vehicle 200 may include first vehicle side 802, which may be first distance 805 from a first lane constraint 2302. Similarly, vehicle 200 may include second vehicle side 812 opposite from first vehicle side 802, and second vehicle side 812 may be second distance 815 from a second lane constraint 2304. In this manner, first lane constraint 810 and second lane constraint 2304 may define a lane 2330 within which vehicle 200 may travel.

[0252] System 100 may determine first lane constraint 2302 and second lane constraint 2304 based on a plurality of images acquired by one or more of image capture devices 122-126. According to some embodiments, first lane constraint 2302 and/or second lane constraint 2304 may be identified by visible lane boundaries, including lines marked on a road surface. Additionally or alternatively, first lane constraint 2302 and/or second lane constraint 2304 may include an edge of a road surface. According to some embodiments, system 100 may determine first lane constraint 2302 and/or second lane constraint 2304 by identifying a midpoint of a road surface width 2320 (or any other suitable road feature that vehicle 200 may use as a navigation reference). System 100 may identify lane constraints in this manner when, for example, lines designating road lanes are not painted or otherwise labeled.

[0253] When passing a target vehicle 2310, the lane change of vehicle 200 may be complicated by a change in the trajectory of target vehicle 2330. Thus, system 100 may identify target vehicle 2310 based on an analysis of a plurality of images acquired by one or more of image capture devices 122-126. System 100 may also identify second lane 2340 in which target vehicle 2310 is traveling. Further, system 100 may determine whether second lane 2340 is different than lane 2330 in which vehicle 200 is traveling. If second lane 2340 is different than lane 2330, system 100 may be configured to enable vehicle 200 to pass target vehicle 2310.

[0254] System 100 may also monitor target vehicle 2310 based on the plurality of images. For example, system 100 may monitor a position of target vehicle 2310, such as monitoring its position relative to vehicle 200, relative to first lane constraint 2302 and/or second lane constraint 2304, an absolute position, or relative to another

Tesla EX1002 Page 1646 IPR2025-00943

reference point. Additionally or alternatively, monitoring a position of target vehicle 2310 may include estimating a speed of target vehicle 2310.

[0255] If system 100 determines that target vehicle 2310 is entering lane 2330 in which vehicle 200 is traveling, system 100 may cause vehicle 200 to abort the pass of target vehicle 2310. For example, system 100 may cause vehicle 200 to stop accelerating or brake. Additionally or alternatively, system 100 may issue an audible announcement, such as to notify driver of vehicle 200 that the lane change is being aborted.

[0256] FIG. 24A is an exemplary block diagram of memory 140, which may store instructions for performing one or more operations consistent with disclosed embodiments. As illustrated in FIG. 24A , memory 140 may store one or more modules for identification of the target vehicle and responses described herein. For example, memory 140 may store a lane constraint module 2400, a target vehicle acquisition module 2410, and an action module 2415.

[0257] In one embodiment, lane constraint module 2400 may store instructions which, when executed by processing unit 110, may detect and define first lane constraint 2302 and second lane constraint 2304. For example, processing unit 110 may execute lane offset module 2400 to process the plurality of images received from at least one image capture device 122-124 and detect first lane constraint 2302 and second lane constraint 2304. As discussed above, this may include identifying painted lane lines or markers and/or measuring a midpoint of a road surface. Processing unit 110 may perform the analysis based on the techniques described in connection with FIGS. 5A-5D , above.

[0258] In some embodiments, target acquisition module 2410 may store instructions which, when executed by processing unit 110, may detect the presence of target vehicle 2310 and monitor the position of target vehicle 2310. For example, target acquisition module 2410 may process the plurality of images to detect and monitor target vehicle 2310. Additionally or alternatively, when executing target acquisition module 2410, processing unit 110 may receive information from another module or other system indicative of the presence and/or location of target vehicle 2310. Processing unit 110 may perform the analysis based on the techniques described in connection with FIGS. 5A-5D , above.

[0259] Processing unit 110 may be configured to determine the presence of one or more objects such as target vehicle 2310 in the vicinity of vehicle 200. For example target vehicle 2310 may comprise another vehicle, such as a car, truck or motorcycle traveling near vehicle 200. Processing unit 110 may determine, for each target vehicle 2310, an offset profile. The offset profile may include a determination of whether in its current and predicted position, target vehicle 2310 will be within a predefined range of vehicle 200. If processing unit 110 determines that target vehicle 2310 is or will be within a predefined range of vehicle 200, processing unit may determine whether there is enough space within first lane constraint 810 and second lane constraint 820, for vehicle 200 to bypass target vehicle 2310. If there is not enough space, processing unit 110 may execute a lane change. If there is enough space, processing unit 110 may determine whether there is enough time to bypass target vehicle 2310 before the distance between vehicle 200 and target vehicle 2310 closes. If there is enough time, processing unit 110 may initiate and or activate the offset maneuver to bypass target vehicle 2310. The offset maneuver may be executed by processing unit 110 so that the movement of vehicle 200 is smooth and comfortable for the driver. For example, the

IPR2025-00943
Tesla EX1002 Page 1647

offset slew rate may be approximately 0.15 to 0.75 m/sec. The offset maneuver may be executed so that the maximum amplitude of the offset maneuver will occur when the gap between vehicle 200 and object 825 closes. If there is not enough time, processing unit 110 may initiate braking and then initiate an offset maneuver.

[0260] Further, processing unit 110 may execute action module 2415 to abort the plan of vehicle 200 to pass target vehicle 2310. If target vehicle 2310 is not in lane 2330 of vehicle 200, instructions included in action module 1230 may enable vehicle 200 to pass target vehicle 2330. However, if execution of target vehicle acquisition module 2410 indicates that target vehicle 2310 is changing lanes into lane 2330 of vehicle 200, processing unit 110 may execute instructions included in action module 2415 to abort the lane change of vehicle 200. For example, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to abort the lane change. Additionally or alternatively, processing unit may cause an audible notification to sound.

[0261] FIG. 24B illustrates a flowchart of an exemplary process 2420 for navigating a vehicle among encroaching vehicles, consistent with disclosed embodiments. Process 2420 may identify lane constraints that define a lane of travel for a vehicle (e.g., vehicle 200) to travel, identify and monitor a target vehicle (e.g., vehicle 2310), and enable the vehicle to pass the target vehicle or abort the pass if it is determined that the target vehicle is entering a lane in which the vehicle is traveling.

[0262] At step 2430, at least one of image capture device 122, 124, and/or 126 may acquire a plurality of images of an area in the vicinity of vehicle 200. For example, a camera included in image acquisition unit 120 (such as image capture device 122) may capture a plurality of images and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive more than one plurality of images via a plurality of data interfaces. For example, processing unit 110 may receive a plurality of images from each of image capture devices 122, 124, 126, each of which may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0263] Next, at step 2440, processing unit 110 may determine a first lane constraint on first vehicle side 802 and a second lane constraint on second vehicle side 812 using the plurality of images received via data interface 128. For example, as part of implementing step 2440, processing unit 110 may execute instructions of lane constraint module 910. Further, as part of determining the first and second lane constraints, processing unit 110 may use one or more of the processes discussed above in connection with FIGS. 5A-5D .

[0264] At step 2450, processing unit 110 may acquire target vehicle 2310. To acquire target vehicle 2310, processing unit 110 may execute target vehicle acquisition module 2410. For example, processing unit 110 may execute instructions to analyze acquired images to identify second lane 2340 in which target vehicle 2310 is traveling and determine whether target vehicle 2310 is traveling in lane 2330 in which vehicle 200 is traveling. This determination may be based on the techniques described in connection with FIGS. 5A-5D , above.

[0265] At step 2455, processing unit 110 may determine whether second lane 2340 in which target vehicle 2310 is traveling is different than lane 2330 in which vehicle 200 is

IPR2025-00943
Tesla EX1002 Page 1648

traveling.

[0266] At step 2460, processing unit 110 may enable vehicle 200 to pass target vehicle 2310 if second lane 2340 is different than lane 2330 in which vehicle 200 is traveling. To enable vehicle 200 to pass target vehicle 2310, processing unit may execute action module 2415.

[0267] Next, at step 2465, processing unit 110 may execute target acquisition module 2410 to monitor target vehicle 2310. This may include estimating a velocity of target vehicle 2310 and determining the location of target vehicle 2310, such as with respect to vehicle 200.

[0268] Step 2470 may include determining whether target vehicle 2310 has entered lane 2330. If processing unit 110 determines that target vehicle 2310 has not entered lane 2330, at step 2480, processing unit 110 may allow vehicle 200 to complete the pass.

[0269] At step 2490, if processing unit 110 determines that target vehicle 2310 has entered lane 2330 (or is entering lane 2330 or is otherwise on a trajectory to bring vehicle 2310 into lane 2330), processing unit 110 may execute action module 2415 to abort the pass. This may include, for example, issuing an audible notification to driver of vehicle 200 via speakers 360. Additionally or alternatively, step 2490 may include executing instructions included in action module 2415 to cause processing unit 110 to transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to abort lane change.

Navigating a Vehicle to Avoid Encroaching Traffic

[0270] System 100 may provide driver assist functionality that monitors the vicinity around vehicle 200 and responds to the presence of another vehicle encroaching upon vehicle 200. Monitoring the vicinity around vehicle 200 may include monitoring a collision threshold. For example a collision threshold may include a time to collision or a minimum predetermined distance between vehicle 200 and other traffic. System 100 may further cause vehicle 200 to take evasive action to avoid encroaching traffic. For example, evasive action may include braking, changing lanes, or otherwise altering the course of vehicle 200.

[0271] As illustrated in FIG. 25 , vehicle 200 may travel on roadway 2500. Vehicle 200 may include system 100 to detect and respond to encroaching traffic. For example, vehicle 200 may determine that another vehicle is entering into the lane in which vehicle 200 is traveling or otherwise determine that another vehicle is crossing a collision threshold. Vehicle 200 may react to an encroaching vehicle, such as by altering course, braking, or accelerating.

[0272] As discussed with respect to FIG. 8 and as shown in FIG. 25 , vehicle 200 may include first vehicle side 802, which may be first distance 805 from a first lane constraint 2502. Similarly, vehicle 200 may include second vehicle side 812 opposite from first vehicle side 802, and second vehicle side 812 may be second distance 815 from a second lane constraint 2504. In this manner, first lane constraint 2502 and second lane constraint 2504 may define a lane 2505 within which vehicle 200 may travel.

Tesla EX1002 Page 1649 IPR2025-00943

[0273] System 100 may determine first lane constraint 2502 and second lane constraint 2504 based on a plurality of images acquired by one or more of image capture devices 122-126. According to some embodiments, first lane constraint 2502 and/or second lane constraint 2504 may be identified by visible lane boundaries, such a line marked on a road surface. Additionally or alternatively, first lane constraint 2502 and/or second lane constraint 2504 may include an edge of a road surface. According to some embodiments, system 100 may determine first lane constraint 2502 and/or second lane constraint 2502 by identifying a midpoint 2520 of a road surface width (or any other suitable road feature vehicle 200 may use as a navigation reference). System 100 may identify lane constraints in this manner when, for example, lines designating road lanes are not painted or otherwise labeled.

[0274] Detection of first lane constraint 2502 and/or second lane constraint 2504 may include processing unit 110 determining their 3D models in a camera coordinate system. For example, the 3D models of first lane constraint 2502 and/or second lane constraint 2504 may be described by a third-degree polynomial. In addition to 3D modeling of first lane constraint 2502 and/or second lane constraint 2504, processing unit 110 may estimate motion parameters, such as the speed, yaw and pitch rates, and acceleration of vehicle 200. Optionally, processing unit may detect static and moving vehicles and their position, heading, speed, and acceleration, all relative to vehicle 200. Processing unit 110 may determine a road elevation model to transform the information acquired from the plurality of images into 3D space.

[0275] The driver of vehicle 200 cannot assume that other traffic traveling in adjacent or near-adjacent lanes will only change lanes if it is safe for that traffic to shift into lane 2505 in which vehicle 200 is driving. Thus, it may be desirable to consider and monitor traffic traveling in nearby lanes to detect lateral encroachment into lane 2505 of vehicle 200, so that vehicle 200 may react to lateral encroachment to avoid a collision. System 100 may identify an encroaching vehicle 2510 based on the plurality of images and determine that encroaching vehicle 2510 is approaching vehicle 200. System 100 may then cause an action to account for the encroaching vehicle. For example, system 100 may cause vehicle 200 to maintain a current velocity and travel within first lane constraint 2502 and second lane constraint 2504 such that a first distance 2530, which is on first vehicle side 802 on which encroaching vehicle 2510 is approaching, is greater than a second distance 2540.

[0276] Processing unit 110 may be configured to determine the presence of one or more encroaching vehicles 2510 in the vicinity of vehicle 200. Processing unit 110 may determine, for each encroaching vehicle 2510 an offset profile. The offset profile may include a determination of whether in its current and predicted position, encroaching vehicle 2510 will be within a predefined range of vehicle 200. If processing unit 110 determines that encroaching vehicle 2510 is or will be within a predefined range of vehicle 200, processing unit may determine whether there is enough space within first lane constraint 2502 and second lane constraint 2504, for vehicle 200 to bypass encroaching vehicle 2510. If there is not enough space, processing unit 110 may execute a lane change. If there is enough space, processing unit 110 may determine whether there is enough time to bypass encroaching vehicle 2510 before the distance between vehicle 200 and encroaching vehicle 2510 closes.

[0277] For example, if there is enough time, processing unit 110 may initiate and or activate the offset maneuver to bypass encroaching vehicle 2510. The offset maneuver may be executed by processing unit 110 so that the movement of vehicle 200 is

IPR2025-00943
Tesla EX1002 Page 1650

smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 to 0.75 m/sec. The offset maneuver may be executed so that the maximum amplitude of the offset maneuver will occur when the gap between vehicle 200 and encroaching vehicle 2510 closes. If there is not enough time, processing unit 110 may initiate braking (e.g., via transmitting electronic signals to braking system 230) and then initiate offset maneuver (e.g., by transmitting electronic signals to throttling system 220 and/or steering system 240).

[0278] System 100 may determine, based on a plurality of images acquired by one or more of image capture devices 122-126, that encroaching vehicle 2510 has crossed at least one collision threshold. For example, the at least one collision threshold may include a minimum predetermined distance 2550 between vehicle 200 and encroaching vehicle 2510. For example, encroaching vehicle 2510, as illustrated in FIG. 25 , is shown exceeding minimum predetermined distance 2550. Additionally or alternatively, the at least one collision threshold may include a time to collision of vehicle 200 and encroaching vehicle 2510. The time-to-collision threshold may be determined based on the distance between vehicle 200 and encroaching vehicle 2510, the velocity of vehicle 200 and/or encroaching vehicle 2510, the lateral velocity of vehicle 200 and/or encroaching vehicle 2510, and/or the acceleration of vehicle 200 and/or encroaching vehicle 2510. A collision threshold based on a time to collision may be set to allow enough time for system 100 to effectively conduct evasive action to avoid collision.

[0279] If system 100 determines that at least one collision threshold has been crossed, system 100 may cause vehicle 200 to take evasive action. For example, this may include causing other subsystems of vehicle 200 to operate. Thus, system 100 may operate to cause vehicle 200 to increase speed by communicating with throttling system 220, braking system 230, and/or steering system 240 to change the speed and/or direction of vehicle 200.

[0280] The evasive action may include causing vehicle 200 to change lanes. This may be a desirable response if, for example, vehicle 200 is traveling on a multilane highway and system 100 determines that changing lanes to move away from encroaching vehicle 2510 will avoid a collision. Evasive action may include vehicle 200 altering its course. Additionally or alternatively, evasive action may include altering course of vehicle 200. For example, to take evasive action, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240.

[0281] System 100 may also be configured to cause vehicle 200 to take an action after the evasive action has been completed. For example, system 100 may cause vehicle 200 to return to midpoint 2520 of lane 2505 in which vehicle 200 is traveling. Additionally or alternatively, system 100 may cause vehicle 200 to resume a preset speed after braking in response to encroaching vehicle 2510.

[0282] FIG. 26 is an exemplary block diagram of memory 140, which may store instructions for detecting and responding to traffic laterally encroaching on a vehicle consistent with disclosed embodiments. As illustrated in FIG. 26 , memory 140 may store a lane constraint module 2610, an encroaching vehicle acquisition module 2620, and an action module 2630.

[0283] In one embodiment, lane constraint module 2610 may store instructions which, when executed by processing unit 110, may detect and define first lane constraint 2502

and second lane constraint 2504. For example, processing unit 110 may execute lane offset constraint module 2610 to process the plurality of images received from at least one image capture device 122-126 and detect first lane constraint 2502 and second lane constraint 2504. As discussed above, this may include identifying painted lane lines and/or measuring a midpoint of a road surface. Processing unit 110 may perform the analysis based on the techniques described in connection with FIGS. 5A-5D , above.

[0284] In one embodiment, encroaching vehicle acquisition module 2620 may store instructions which, when executed by processing unit 110, may detect the presence of encroaching vehicle 2510 and monitor encroaching vehicle 2510. For example, encroaching vehicle acquisition module 2620 may include instructions for determining a relative location of encroaching vehicle 2510 to vehicle 200. In some embodiments, this may include monitoring the velocity of encroaching vehicle 200. According to some embodiments, encroaching vehicle acquisition module 2620 may determine whether encroaching vehicle 2510 has exceeded a collision threshold. A collision threshold may be, for example, a minimum predetermined distance between vehicle 200 and encroaching vehicle 2510, If encroaching vehicle 2510 comes closer to vehicle 200 than the minimum predetermined distance, processing unit 110 may determine that encroaching vehicle 2510 is indeed encroaching and execute instructions to avoid collision with encroaching vehicle 2510.

[0285] Additionally or alternatively, a collision threshold may include a time to collision. The time-to-collision threshold may be determined based on the distance between vehicle 200 and encroaching vehicle 2510, the velocity of vehicle 200 and/or encroaching vehicle 2510, the lateral velocity of vehicle 200 and/or encroaching vehicle 2510, and/or the acceleration of vehicle 200 and/or encroaching vehicle 2510. A collision threshold based on a time to collision may be set to allow enough time for processing unit 110 and system 100 to effectively conduct evasive action to avoid collision.

[0286] Further, processing unit 110 may execute action module 2630 to respond to the detection of encroaching vehicle 2510 and/or the crossing of a collision threshold. For example, action module 2630 may include instructions to cause vehicle 200 to maintain a current velocity and to travel within first lane constraint 2502 and second lane constraint 2504 such that first distance 2530 is greater than second distance 2540. The first distance 2530 may be on the same side of vehicle 200 as encroaching vehicle 2510, so that vehicle 200 is farther away from encroaching vehicle 2510 while still within first lane constraint 2502 and second lane constraint 2504. If processing unit 110 determines that a collision threshold has been crossed, processing unit 110 may execute action module 2630 to cause vehicle 200 to take evasive action, such as altering course, braking, and/or changing lanes. Action module 2630 may also include instructions that cause vehicle 200 to resume course after completing evasive action. For example, processing unit 110 may cause vehicle 200 to resume a preset speed after braking or return to the center of lane 2505.

[0287] FIG. 27 illustrates a flowchart of an exemplary process 2720 detecting and responding to traffic laterally encroaching on a vehicle, consistent with disclosed embodiments. Process 2720 may identify lane constraints that define a lane of travel for vehicle travel, determine whether an encroaching vehicle 2510 is approaching, and cause vehicle 200 to maintain current velocity and to travel within first lane constraint 2501 and second lane constraint 2504 such that first distance 2530, on the side of

vehicle 200 that encroaching vehicle 2510 is approaching, is greater than second distance 2540.

[0288] At step 2730, at least one image capture device 122,124, and/or 126 may acquire a plurality of images of an area in the vicinity of vehicle 200. For example, a camera included in image acquisition unit 120 (such as image capture device 122) may capture a plurality of images and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive more than one plurality of images via a plurality of data interfaces. For example, processing unit 110 may receive a plurality of images from each of image capture devices 122, 124, 126, each of which may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0289] Next, at step 2740, processing unit 110 may determine a first lane constraint on first vehicle side 802 and a second lane constraint on second vehicle side 812 using the plurality of images received via data interface 128. For example, as part of implementing step 2740, processing unit 110 may execute instructions of lane constraint module 2610. Further, as part of determining the first and second lane constraints, processing unit 110 may use one or more of the processes discussed above in connection with FIGS. 5A-5D .

[0290] At step 2750, processing unit 110 may determine whether encroaching vehicle 2510 is approaching. To make this determination, processing unit 110 may execute encroaching vehicle acquisition module 2620. Encroaching vehicle acquisition module 2620 may include instructions to process images to detect other traffic. This may include identifying vehicles within a certain range of vehicle 200, such as those vehicles that are adjacent to vehicle 200. These adjacent vehicles may be encroaching vehicles 2510.

[0291] Once encroaching vehicle 2510 is identified, processing unit may determine more information regarding vehicle 2510. For example, at step 2555, processing unit 110 may determine whether target vehicle 2330 has exceeded a collision threshold. A collision threshold may be, for example, a minimum predetermined distance between vehicle 200 and encroaching vehicle 2510. If encroaching vehicle 2510 comes closer to vehicle 200 than the minimum predetermined distance, processing unit 110 may determine that encroaching vehicle 2510 is indeed encroaching and execute instructions to avoid collision with encroaching vehicle 2510.

[0292] Additionally or alternatively, a collision threshold may be a time to collision. The time-to-collision threshold may be determined based on the distance between vehicle 200 and encroaching vehicle 2510, the velocity of vehicle 200 and/or encroaching vehicle 2510, the lateral velocity of vehicle 200 and/or encroaching vehicle 2510, and/or the acceleration of vehicle 200 and/or encroaching vehicle 2510. A collision threshold based on a time to collision may be set to allow enough time for processing unit 110 and system 100 to effectively conduct evasive action to avoid collision.

[0293] At step 2760, processing unit 110 may cause vehicle 200 to maintain current velocity and to travel within first lane constraint 2502 and second lane constraint 2504. To cause vehicle 200 to operate in this manner, processing unit 110 may execute action module 2630. Additionally or alternatively, step 2760 may include executing instructions on action module 1230 to cause processing unit 110 to transmit electronic

Tesla EX1002 Page 1653 IPR2025-00943

signals to the accelerator 2610, brakes 2620, and/or steering system of vehicle 200 to cause vehicle 200 to maintain current velocity and to travel within first and second lane constraints 810 and 820.

[0294] If, at step 2755, processing unit 110 determined that a collision threshold has been exceeded, the processing unit 110 may take evasive action. The evasive action may include causing vehicle 200 to change lanes. This may be a desirable response if, for example, vehicle 200 is traveling on a multilane highway and system 100 determines that changing lanes to move away from encroaching vehicle 2510 will avoid a collision. Evasive action may include vehicle 200 altering its course. Additionally or alternatively, evasive action may include altering course of vehicle 200. For example, processing unit 110 may undertake evasive action by transmitting electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. After the evasive action is completed, at step 2780, processing unit 110 may control vehicle 200 to resume its prior navigation (e.g., to resume navigation to a desired location).

Multi-Threshold Reaction Zone for Vehicle Navigation

[0295] System 100 may provide driver assist functionality that causes a response in vehicle 200, such as braking, accelerating, switching lanes, turning, and/or other navigational responses. For example, if system 100 detects the presence of another vehicle ahead of vehicle 200 that is traveling at a rate of speed slower than vehicle 200, system 100 may cause vehicle 200 to reduce its speed by braking. System 100 may apply a braking profile that accounts for the immediacy of the risk (e.g., the risk of collision) to provide the driver with a natural driving sensation.

[0296] System 100 may cause various different responses depending on the situation and a determined risk of collision with another vehicle, for example. System 100 may take no action in response to changes in speed or direction of a target vehicle ahead if the target vehicle is sufficiently far ahead and the determined risk of collision is low. Where the target vehicle may be closer or when the risk of collision rises above a predetermined level, system 100 may cause a response, such as braking, etc. When the risk of collision is even higher (e.g., where system 100 determines that a collision would be imminent without evasive action), system 100 may cause vehicle 100 to take an evasive action, such as maximum or near maximum braking, change of direction, etc.

[0297] FIGS. 28A and 28B are diagrammatic representations of a primary vehicle 200a on a road 2800, consistent with disclosed embodiments. In some instances, road 2800 may be divided into two lanes, lane 2810 and lane 2820. A target object, such as target vehicle 200b, may be traveling ahead of primary vehicle 200a in the same lane (e.g., lane 2820 as shown in FIGS. 28A and 28B) or in a different lane. As shown in the example depicted in FIG. 28A , primary vehicle 200a and target vehicle 200b may be separated by a distance d1. At a given time after the scene depicted in FIG. 28A , primary vehicle 200a and target vehicle 200b may be separated by a distance d2as shown in FIG. 28B , where d2is less than d1. For example, primary vehicle 200a may be traveling at a rate of speed greater than target vehicle 200b, causing the distance between the vehicles to decrease from d1to d2.

[0298] As another example, primary vehicle 200a and target vehicle 200b may be traveling at the same rate of speed until target vehicle 200b reduces its speed by

Tesla EX1002 Page 1654 IPR2025-00943

braking, or primary vehicle 200a increases its speed by accelerating, causing the distance between the vehicles to decrease from d1to d2. As another example, target vehicle 200b may encroach from a side of primary vehicle 200a. In some circumstances, if target vehicle 200b is ahead or a sufficient distance away from primary vehicle 200a (e.g., five meters or more, 10 meters or more, etc.) and moves toward primary vehicle 200a, system 100 may not take any action. However, if target vehicle 200b continues to move closer to primary vehicle 200a, then system 100 may cause primary vehicle 200a to brake, change speed, and/or change lanes. For example, if target vehicle 200b is within a predetermined threshold (e.g., within five meters), primary vehicle 200a may then take action. Different scenarios involving the relative speeds and positions of primary vehicle 200a and target vehicle 200b are possible, and the disclosed embodiments are not limited to any particular scenario.

[0299] In some embodiments, the multi-threshold reaction zone may provide a natural feeling driving experience to the user. For example, as discussed above, if the risk of collision is low, then system 100 may disregard movements by target vehicle 200b. As a result, primary vehicle 200a may not respond to movements or actions of target vehicle 200b when there is not a risk of a collision (e.g., primary vehicle 200a would not need to mimic every motion of target vehicle 200b when it's 100 meters ahead). However, if target vehicle 200b is close to primary vehicle 200a and/or if system 100 determines that the risk of collision is high, then a responsive maneuver may be taken such that the user feels safer by increasing space between primary vehicle 200a and target vehicle 200b. If system 100 determines that a crash may be imminent, then system 100 may take a more dramatic action without causing the user to be surprised, because the action may be necessary to avoid a collision.

[0300] Accordingly, in some embodiments, system 100 may associate the different reaction zones with different degrees of responsive actions. For example, in some embodiments, system 100 may take no action when target vehicle 200b is within a first threshold (e.g., when target vehicle 200b is far away from primary vehicle 200a), system 100 may take a medium action with target vehicle 200b is within a second threshold (e.g., when target vehicle 200b is becoming closer to primary vehicle 200a), and system 100 may take a more dramatic action when target vehicle 200 is within a third threshold (e.g., when target vehicle 200b is sufficiently close to primary vehicle 200a that there is a risk of a collision). In a first reaction zone associated with the first threshold, system 100 may take no action when the estimated distance-to-target associated with target vehicle 200b is greater than twice a predetermined safety distance. In some embodiments, the first reaction zone may apply where the estimated distance-to-target associated with target vehicle 200b falls in the range of 1.5 to 2.5 times the predetermined safety distance.

[0301] The estimated distance-to-target may be the current distance between primary vehicle 200a and target vehicle 200b plus the expected change of this distance over a prediction time (e.g., in the range of 0.1 to 1.0 seconds, depending on factors associated with system 100 such as velocity tracking control loops, actuators, and vehicle dynamics). The expected change of distance may be calculated as: (vtarget-vprimary)*tprediction+ ((atarget-aprimary)*tprediction<1/2>), wherevtargetandvprimarycorrespond to the speed of target vehicle 200b and primary vehicle 200a, respectively,atargetandaprimarycorrespond to the acceleration of target vehicle 200b and primary vehicle 200a, respectively, andtpredictioncorresponds to the prediction time. Choosing a prediction time approximating the overall reaction time of system 100 may account for the fact that desired navigational responses for primary

vehicle 200a may not be implemented immediately, and thereby may provide the driver and/or passengers of primary vehicle 200a with a smooth, natural driving experience.

[0302] The predetermined safety distance may be defined as: max {dstatic , ddynamic}.dstatic ,which may fall in the range of 2 to 10 meters, may represent a desired safety distance between primary vehicle 200a and another vehicle (such as target vehicle 200b) while the vehicles are stopping and/or moving at a low rate of speed.ddynamicmay represent a desired safety distance when the vehicles are moving at a speed greater than a low rate of speed, and may be calculated as:tsafety*min (vtarget , vprimary), wherevtargetandvprimarycorrespond to the speed of target vehicle 200b and primary vehicle 200a, respectively.tsafetymay fall in the range of 0.5 to 2.5 seconds and may be adjusted by the driver of primary vehicle 200a via user interface 170 according to the driver's preferences. Thus, the driver may be able to control the distance that system 100 maintains between primary vehicle 200a and other vehicles.

[0303] In a second reaction zone associated with the second threshold, system 100 may place target vehicle 200b in a safety zone and apply different weights to different maneuvers performed by target vehicle 200b. For example, as target vehicle 200b may approach a lower boundary of the safety zone, e.g., the estimated distance-to-target associated with target vehicle 200b may approach the predetermined safety distance (discussed above in connection with the first reaction zone), system 100 may assign more weight to deceleration maneuvers (e.g., mimicked more) and less weight to acceleration maneuvers (e.g., mimicked less). In some embodiments, the lower boundary may fall in the range of 0.5 to 1.5 times the predetermined safety distance, depending on driver preferences. Conversely, as target vehicle 200b approaches an upper boundary of the safety zone, e.g., the estimated distance-to-target associated with target vehicle 200b may approach twice the predetermined safety distance, system 100 may assign less weight to deceleration maneuvers (e.g., mimicked less) and more weight to acceleration maneuvers (e.g., mimicked more). In some embodiments, the lower boundary may fall in the range of 1.5 to 2.5 times the predetermined safety distance, depending on driver preferences.

[0304] In a third reaction zone associated with the third threshold, system 100 may place target vehicle 200b in a danger zone when the estimated distance-to-target associated with target vehicle 200b is less than the predetermined safety distance. In this scenario, system 100 may cause primary vehicle 200a to mimic (e.g., immediately) any deceleration maneuvers (while disregarding any acceleration maneuvers) performed by target vehicle 200b. In some embodiments, the first reaction zone may apply where the estimated distance-to-target associated with target vehicle 200b falls in the range of 0.5 to 1.5 times the predetermined safety distance, depending on driver preferences.

[0305] In some embodiments, system 100 may cause a response in primary vehicle 200a, such as braking. As indicated in FIG. 28B , for example, the distance between the primary vehicle 200a and target vehicle 200b may decrease from d1to d2. Depending on the distance d2and, by extension, the time before collision, system 100 may cause primary vehicle 200a to reduce its speed so as to avoid a collision with target vehicle 200b. In addition, system 100 may apply one or more braking profiles when causing primary vehicle 200a to reduce its speed, also based on, for example, the time before collision. Thus, if the time before collision (and/or the distance d2) exceeds a first predetermined threshold, system 100 may cause primary vehicle 200a to brake in a gradual manner, change lanes, and/or change speed, whereas if the time

Tesla EX1002 Page 1656 IPR2025-00943

before collision (and/or the distance d2) is below a second predetermined threshold, system 100 may cause primary vehicle 200a to brake, change lanes, and/or change speed in a more rapid manner.

[0306] FIG. 29 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations, consistent with the disclosed embodiments. As shown in FIG. 29 , memory 140 may store a target monitoring module 2902, an intercept time module 2904, and an action response module 2906. The disclosed embodiments are not limited to any particular configuration of memory 140. Further, processing unit 110 may execute the instructions stored in any of modules 2902-2906 included in memory 140.

[0307] In one embodiment, target monitoring module 2902 may store software instructions (such as computer vision software) which, when executed by processing unit 110, identifies and monitors a target object (such as target vehicle 200b) appearing in a plurality of images representing an environment in the vicinity of primary vehicle 200. For purposes of this disclosure, the target object may also be referred to as target vehicle 200b, although the target object may be a pedestrian, road hazard (e.g., large debris), and the like. The plurality of images may be acquired by one or more of image capture devices 122, 124, and 126. Processing unit 110 may execute target monitoring module 2902 to identify and monitor target vehicle 200b by, for example, analyzing the plurality of images using the techniques described in connection with FIGS. 5A-5D , above. Thus, target monitoring module 2902 may track different variables, such as the motion (e.g., direction of travel and trajectory), speed, and acceleration associated with target vehicle 200b, as well as the distance between primary vehicle 200a and the target vehicle 200b. Such information may be tracked continuously or at predetermined intervals.

[0308] In one embodiment, intercept time module 2904 may store software instructions which, when executed by processing unit 110, determines an indicator of an intercept time between primary vehicle 200a and target vehicle 200b. The intercept time may refer to an amount of time before primary vehicle 200a makes contact (e.g., collides) with target vehicle 200b. Processing unit 110 may execute intercept time module 2904 to determine an indicator of the intercept time based on an analysis of images acquired by one or more of image capture devices 122, 124, and 126. For example, intercept time module 2904 may make the determination using a series of time-based observations, e.g., of the position, speed, and/or acceleration of target vehicle 200b (relative to primary vehicle 200a), such as Kalman filters or linear quadratic estimation (LQE). The Kalman filters may be based on a measurement of the scale of target vehicle 200b, where the scale measurement is proportional to the intercept time.

[0309] In one embodiment, action response module 2906 may store software instructions (such as computer vision software) which, when executed by processing unit 110, causes one or more responses in primary vehicle 200a. The responses may be based on the identification and monitoring of target vehicle 200b (e.g., performed via execution of target monitoring module 2902) and/or the indicator of an intercept time between primary vehicle 200a and target vehicle 200b (e.g., determined via execution of intercept time module 2904). Processing unit 110 may cause one or more responses in primary vehicle 200a by, for example, transmitting electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to brake, accelerate, or trigger a turn, lane shift, and/or change in direction of primary vehicle 200a. Processing unit 110 may use extended Kalman

Tesla EX1002 Page 1657 IPR2025-00943

Filters to estimate data associated with target vehicle 200b and primary vehicle 200a, such as: speed, acceleration, turn rate, pitch angle, pitch rate, position and heading of target vehicle 200b relative to primary vehicle 200a, the size of target vehicle 200b, road geometry, and the like. For example, processing unit 110 may receive measurements from components on primary vehicle 200a, such as an odometer, an inertial measurement unit (IMU), image processor 190, and the like. Thus, processing unit 110 may calculate data associated with target vehicle 200b and/or primary vehicle 200a using state equations derived from kinematic equations of motion, rigidity constraints, and motion assumptions, where inputs to the equations may include measurements from components on primary vehicle 200a. Processing unit 110 may also use the techniques described in connection with FIGS. 4-7 , above, to cause one or more responses in primary vehicle 200a.

[0310] FIG. 30 is a flow chart showing an exemplary process 3000 for causing a response in a primary vehicle, consistent with disclosed embodiments. At step 3010, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120. For instance, one or more cameras included in image acquisition unit 120 may capture a plurality of images of an area forward of primary vehicle 200a (or to the sides or rear of a vehicle, for example) and transmit them over a data connection (e.g., digital, wired, USB, wireless, Bluetooth, etc.) to processing unit 110.

[0311] At step 3020, processing unit 110 may execute target monitoring module 2902 to identify a target object (such as target vehicle 200b) within the plurality of images. For example, target monitoring module 2902 may detect the presence of vehicles in the images using the techniques described in connection with FIGS. 5A and 5B, above. Multiple vehicles (and/or objects) may be detected in the same lane or in a different lane as primary vehicle 200a, and at varying distances from primary vehicle 200a, based on image data, position data (e.g., GPS location information), map data, the yaw rate of primary vehicle 200a, lane or other road markings, speed data, and/or data from sensors included in primary vehicle 200a. In a scenario where multiple vehicles are detected, target monitoring module 2902 may determine target vehicle 200b to be the vehicle traveling in the same lane as primary vehicle 200a and/or located closest to primary vehicle 200a.

[0312] At step 3030, processing unit 110 may execute target monitoring module 2902 to monitor target vehicle 200b identified at step 3020. For example, target monitoring module 2902 may track information associated with target vehicle 200b, such as the motion (e.g., direction of travel and trajectory), speed, and acceleration, as well as the distance between primary vehicle 200a and the target vehicle 200b. Target monitoring module 2902 may track such information continuously, or at predetermined intervals, by analyzing the plurality of images using the techniques described in connection with FIGS. 5A-5D , above. At step 3040, processing unit 110 may execute intercept time module 2904 to determine an indicator of an intercept time between primary vehicle 200a and target vehicle 200b. Intercept time module 2904 may make the determination by tracking the position, speed, and/or acceleration of target vehicle 200b (relative to primary vehicle 200a) over one or more periods of time. Intercept time module 2904 may use the techniques described in connection with FIGS. 5A and 5B , above to track such information.

[0313] At step 3050, processing unit 110 may execute action response module 2906 to cause one or more responses in primary vehicle 200a. For example, processing unit

IPR2025-00943
Tesla EX1002 Page 1658

110 may transmit electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to brake, accelerate, or trigger a turn, lane shift, and/or change in direction of primary vehicle 200a. Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240. For example, processing unit 110 may transmit electronic signals that cause system 100 to physically depress the brake by a predetermined amount or ease partially off the accelerator of vehicle 200. Further, processing unit 110 may transmit electronic signals that cause system 100 to steer vehicle 200 in a particular direction. Such responses may be based on the monitoring of target vehicle 200b (performed at step 3030) and/or the indicator of an intercept time between primary vehicle 200a and target vehicle 200b (determined at step 3040). In some embodiments, the intercept time may be compared with one or more predetermined thresholds, e.g., a first and second intercept threshold. For example, the response may include an emergency avoidance action if the indicator of intercept time indicates that a target object is within a particular intercept threshold (e.g., within the second intercept threshold). The emergency avoidance action may include changing lanes and/or emergency braking.

[0314] The thresholds may be expressed as units of time. Thus, for example, if the indicator of intercept time indicates that target vehicle 200b is outside of the first threshold (e.g., the time before collision between primary vehicle 200a and target vehicle 200b exceeds the first threshold), then action response module 2906 may cause a response in primary vehicle 200a based on an averaged motion profile associated with target vehicle 200b. The averaged motion profile associated with target vehicle 200b may represent an averaged position and/or averaged speed of target vehicle 200b for a particular threshold. In some embodiments, if target vehicle 200b is a certain distance and/or traveling a certain speed, system 100 may elect to take no action. For example, if target vehicle 200b is ahead or a sufficient distance away from primary vehicle 200a (e.g., five meters, 10 meters, etc.), system 100 may not take any action unless target vehicle 200b becomes closer to primary vehicle 200.

[0315] If the indicator of intercept time indicates that target vehicle 200b is between the first and second thresholds (e.g., the time before collision between primary vehicle 200a and target vehicle 200b falls in between the first and second thresholds), then action response module 2906 may cause a response in primary vehicle 200a based on the actual motion associated with target vehicle 200b. The actual motion may represent a position and/or speed of target vehicle 200b (e.g., not averaged or delayed) If the indicator of intercept time indicates that target vehicle 200b is within the second threshold (e.g., the time before collision between primary vehicle 200a and target vehicle 200b is below the second threshold), then action response module 2906 may cause an emergency action in primary vehicle 200a (such as a lane change or emergency braking).

[0316] FIG. 31 is a flow chart showing an exemplary process 3100 for declining to cause a response in a primary vehicle, consistent with disclosed embodiments. At step 3110, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 3010 of FIG. 30 , above.

[0317] At step 3120, processing unit 110 may execute target monitoring module 2902 to locate a target object (e.g., target vehicle 200b) within the images received at step 3010 by analyzing the images, as described in connection with step 3020 of FIG. 30 ,

IPR2025-00943
Tesla EX1002 Page 1659

above. At step 3125, processing unit 110 may execute target monitoring module 2902 to determine whether the target object (e.g., target vehicle 200b) is traveling in a lane different from primary vehicle 200a, using the techniques described in connection with step 3030 of FIG. 30 , above. If target vehicle 200b is determined to be traveling in a lane different from primary vehicle 200a (step 3125, yes), processing unit 110 may decline to cause a response in primary vehicle 200a at step 3130. Alternatively, if system 100 determines that target vehicle 200b does not constitute an imminent risk of collision (e.g., target vehicle 200b a nonintersecting trajectory with primary vehicle 200a or a potential time to collision beyond a threshold such that the target object does not pose an imminent risk at the present time), processing unit 110 may decline to cause a response in primary vehicle 200a at step 3130. Otherwise (step 3125, no), process 3100 may conclude and no response occurs.

Stopping an Autonomously Driven Vehicle Using a Multi-Segment Braking Profile

[0318] System 100 may provide driver assist functionality that monitors and responds to environmental conditions to control the braking of vehicle 200. For instance, system 100 may detect, based on a plurality of images, an object that triggers the stopping of vehicle 200. For example, this object may include another vehicle in the vicinity of vehicle 200, a traffic light, and/or a traffic sign. System 100 may use a braking profile to slow and/or stop vehicle 200. According to some embodiments, the braking profile may include multiple segments, where deceleration is executed at different stages. For example, the braking profile may cause system 100 to stop vehicle 200 at a decreasing rate of deceleration.

[0319] By providing a braking profile that includes multiple segments, system 100 may cause vehicle 200 to operate in a manner that feels natural to a driver and/or passengers. For example, when a driver stops a vehicle at a traffic light, the brakes are typically not applied at a constant rate from the moment they are applied until the moment when the vehicle stops. Operating brakes in such a manner result in an abrupt stop at the end. Instead, the driver may apply the brakes gently at first, then with increasing pressure, and then as the vehicle reaches its target stopping location, the driver may decrease the level of braking until it reaches zero (or near zero) at the moment the car stops. Operating the vehicle in this manner may provide a natural stopping progression, i.e., the driver may cause-a progressive deceleration at first so the vehicle does not experience an abrupt initial deceleration and then at the end, so that vehicle does not come to an abrupt stop.

[0320] System 100 may provide autonomous vehicle navigation that may implement braking according to an approach that feels natural to the driver and/or passengers of vehicle 200. For example, in some embodiments, system 100 may cause one or more of image capture devices 122-126 to acquire a plurality of images of an area in a vicinity of vehicle 200. Processing device 110 may receive the receive the plurality of images via data interface 128 and identify, based on analysis of the plurality of images, a trigger (e.g., an object, a traffic light, and/or a traffic sign) for stopping vehicle 200. Based on the identified trigger, system 100 may cause vehicle 200 to stop according to a braking profile. In some embodiments, the braking profile may include a plurality of segments (e.g., including a first segment associated with a first deceleration rate, a second segment which includes a second deceleration rate less than the first deceleration rate, and a third segment in which a level of braking is decreased as a target stopping location is approached, as determined based on the analysis of the plurality of images.

IPR2025-00943
Tesla EX1002 Page 1660

[0321] FIG. 32A illustrates vehicle 200 traveling on a roadway 3210 that includes an object 3220. System 100 may identify a trigger (e.g., object 3220) for stopping vehicle 200 in respect to detection of the trigger. For example, system 100 may first slow vehicle 200 after the trigger is detected and bring vehicle 200 to a stop before reaching object 3220.

[0322] Processing unit 110 may be configured to determine a trigger for stopping vehicle 200 based on a plurality of images acquired by one or more of image capture devices 122-126. Processing unit 110 may receive the images from image capture devices 122-126 via data interface 128. According to some embodiments, the trigger may include object 3220. For example, object 3220 may be another vehicle on roadway 2810. According to some embodiments, processing unit 110 may identify a trigger if object 3220 is within a certain vicinity of vehicle 200. For example, identification of the trigger may depend on a relative distance between object 3220 and vehicle 200. Further, this identification may optionally consider the relative velocity and/or relative acceleration between object 3220 and vehicle 200.

[0323] Additionally or alternatively, the trigger may include a traffic sign, such as traffic sign 3240 shown in FIG. 32A . For example, traffic sign 3240 may be a stop sign and/or a speed limit sign. According to some embodiments, processing unit 110 may process the plurality of images to determine the content of traffic sign 3240 as part of determining whether a trigger exists. For example, processing unit 110 may determine that a trigger is present based on a speed limit specified by traffic sign 3240 if vehicle 200 is exceeding the speed limit. As another example, processing unit 110 may determine a trigger exists based on a traffic light 2830 if the yellow and/or red light is on.

[0324] Based on the plurality of images, processing unit 110 may determine a target stopping location 3250 at which vehicle 200 will stop once a braking profile is executed. Reaching target stopping location 3250 may be accompanied by a release in braking pressure of vehicle 200. Target stopping location 3250 may be based on the location of traffic light 3230 and/or traffic sign 3240 such that vehicle 200 complies with traffic light 3230 and/or traffic sign 3240. Alternatively, target stopping location 3250 may be based on object 3220 such that vehicle 200 stops before reaching object 3220. Accordingly, target stopping location 3250 may be based on the location of object 3220 to avoid a collision between vehicle 200 and object 3220. In some embodiments, target stopping location 3250 may be updated based on new information obtained from the plurality of images, such as newly identified targets, newly visible obstacles, newly visible signs, etc.

[0325] FIG. 32B illustrates an exemplary braking profile, consistent with disclosed embodiments. In the example shown in FIG. 32B , the braking profile includes four segments. However, the number of segments is exemplary and a braking profile consistent with the disclosed embodiments may include any appropriate number of segments (e.g., 2, 3, 4, 5, 6, etc., segments).

[0326] In segment 1, system 100 may prepare for a possible braking scenario. For example, in this zone, system 100 may receive information (e.g., map data and/or image data) about an upcoming intersection (e.g., an intersection with traffic lights). System 100 may gradually reduce the speed of vehicle 200 (e.g., the speed of vehicle 200 may be gradually reduced to about 70-50 km/hr, depending on the characteristics

Tesla EX1002 Page 1661 IPR2025-00943

of the intersection). System 100 may cause the speed adjustment, for example, approximately 100-200 meters before the junction at average deceleration around 0.2-0.4 m/sec<2>(with a maximum of about 0.5 m/sec<2>). By decreasing the speed of vehicle 200 in segment 1, system 100 may provide the driver and/or passengers of vehicle 200 with confidence that system 100 is aware of the approaching intersection.

[0327] In segment 2, system 100 may apply strong braking. In this segment, system 100 may receive information based on the analysis of image data that, for example, a traffic light is red and/or there is a stopped vehicle at the intersection in the line in which vehicle 200 is traveling. System 100 may cause vehicle 200 to experience significant speed reduction in this zone, while the distance to the stop line or stopped vehicle is still large. Doing so may provide the driver and/or passengers of vehicle 200 with a comfortable feeling that system 100 will have more than enough space to complete the braking maneuver. This zone may be approximately 30-100 meters away from the stop line, and the average deceleration may be about 1.5-2.5 m/sec<2>(with a maximum of about 3.5 m/sec<2>).

[0328] In segment 3, system 100 may make a moderate braking adjustment. For example, in this segment, system 100 may adjust the speed of vehicle 200 according to the remaining distance to the stop line or stopped vehicle and based on the current speed and deceleration of vehicle 200. By making this adjustment, system 100 may provide the driver and/or passengers of vehicle 200 with an indication that system 100 is releasing most of the braking power while the speed is still being reduced but slower than before. This zone may be approximately 5-30 meters away from the stop line or stopped vehicle, and the average deceleration may be about 0.5-1.5 m/sec<2>(with a maximum of about 2 m/sec<2>).

[0329] In segment 4, system 100 may make a small braking adjustment. In this segment, system 100 may close the remaining distance to the stop line or stopped vehicle at a very low speed gradually bringing it to zero. In this zone, the driver and/or passengers may feel that vehicle 200 is slowly sliding into its position (e.g., at the stop line or behind another vehicle). This zone may constitute approximately the last 5-7 meters to the stop line or stopped vehicle, and the average deceleration may be about 0.3-0.5 m/sec<2>(with a maximum of about 1 m/sec<2>).

[0330] FIG. 33 is an exemplary block diagram of memory 140 and/or 150, which may store instructions for performing one or more operations consistent with disclosed embodiments. As illustrated in FIG. 33 , memory 140 may store one or more modules for performing the trigger detection and responses described herein. For example, memory 140 may store a trigger identification module 3300 and an action response module 3310.

[0331] Trigger identification module 3300 may store instructions which, when executed by processing unit 110, may detect the presence and/or existence of a trigger for stopping vehicle 200. For example, trigger identification module 3300 may process the plurality of images received from at least one image capture device 122-124 to detect the presence of a trigger. As discussed above, this may include identifying object 3220, traffic light 3230, and/or traffic sign 3240.

[0332] Action response module 3310 may store instructions which, when executed by processing unit 110, may respond to the presence of a trigger. For example, action response module 3310 may execute control to decelerate and/or stop vehicle 200. For

Tesla EX1002 Page 1662 IPR2025-00943

example, action response module 3310 may execute a braking profile to decelerate and/or stop vehicle 200. The action response module 3310 may determining to execute one of a plurality of braking profiles based on a number of braking profiles based on, for example, environmental factors, the type of trigger detected, and the velocity of vehicle 200. To navigate vehicle 200 to a stop based on the detected trigger, processing unit 110 may transmit electronic signals to one or more of throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240. For example, processing unit 110 may transmit electronic signals that cause system 100 to physically depress the brake by a predetermined amount.

[0333] Fig. 34 illustrates a process 3400 for navigating vehicle 200, consistent with disclosed embodiments. Process 3000 may identify a trigger (e.g., an object, vehicle, traffic light, traffic sign, etc.) for stopping vehicle 200 and execute a braking profile in response to the detection of the trigger. Process 3400 may detect various types of triggers determine which of a plurality of braking profiles to execute.

[0334] The braking profile may be based on the type of a detected trigger and/or other environmental conditions related to the detected trigger (e.g., a distance to an object). For example, in some embodiments, the braking profile may include a plurality of segments (e.g., three segments). In a three segment braking profile, a first segment of the braking profile may be associated with a first deceleration rate, a second segment of the braking profile may include a second deceleration rate less than the first deceleration rate, and a third segment of the braking profile may include causing a level of braking to be decreased as a target stopping location is approached, as determined based on an analysis of a plurality of images acquired by one or more of image capture devices 122-126. For example, in some embodiments, the first segment may cause vehicle 200 to progressively increase deceleration, the second segment may cause vehicle 200 to decelerate at a constant rate, and the third segment may cause vehicle 200 to progressively decrease its deceleration rate to zero (or near zero).

[0335] At step 3410, process 3400 may acquire, using at least one of image capture devices 122-126, a plurality of images of an area in the vicinity of vehicle 200. For example, processing unit 110 may receive the plurality of images may through data interface 128. Processing unit 110 may be configured to determine the presence of one or more objects and/or signs located on or near roadway 3210 based on the acquired images. For example, processing unit 110 may be configured to determine the presence of object 3220, traffic light 3230, and/or traffic sign 3240. Processing unit 110 may be further configured to read the content of traffic sign 3240. For example, processing unit 110 may be configured to differentiate between a speed limit sign and a stop sign. Further, processing unit 110 may be configured to process images to determine the speed limit based on traffic sign 3240.

[0336] At step 3420, process 300 may include a trigger for stopping vehicle 200. For example, processing unit 110 may identify a trigger based on the presence of object 3220 within the vicinity of vehicle 200. Additionally or alternatively, processing unit 110 may identify a trigger based on a current signal status of traffic light 3230. According to some embodiments, processing unit 110 may determine that traffic light 3230 constitutes a trigger if it is not green. According to some embodiments, processing unit 110 may determine that traffic sign 3240 constitutes a trigger if it is a stop sign, for example. In some embodiments, traffic sign 3240 may constitute a trigger if it is a speed limit sign and vehicle 200 is exceeding the posted speed limit. Further in some

IPR2025-00943
Tesla EX1002 Page 1663

embodiments, processing unit 110 may cause system 100 to provide an audible announcement identifying the trigger and/or indicating that vehicle 200 is braking or about to begin braking.

[0337] After processing unit 110 identifies the trigger, at step 3430, processing unit 110 may cause vehicle 200 to slow down and/or stop according to a braking profile. In some embodiments, the braking profile may comprise a number of profile segments. The different profile segments of a braking profile may be defined by a time period and/or a vehicle speed. For example, processing unit 110 may execute controls of vehicle 200 in accordance with a first profile segment for a predetermined period of time or until the speed of vehicle 200 reaches a target value. The target value may be a predetermined speed or it may be a percentage of the speed that vehicle 200 was traveling prior to execution of the braking profile. To implement the braking profile, processing unit 110 may transmit electronic signals to one or more of throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240.

[0338] As discussed above, according to some embodiments, a profile segment of the braking profile may correspond to a deceleration rate. For example, a braking profile may include multiple segments, and the first segment may be associated with a first deceleration rate. The braking profile may further include a second segment associated with a second deceleration rate. The second deceleration rate may be less than first deceleration rate.

[0339] Further, as discussed above, according to some embodiments, braking profile may include a third profile segment in which a level of braking is decreased as a function of the distance between vehicle 200 and underlying factors that give rise to the trigger (e.g., object 3220, traffic light 3230, traffic sign 3240). For example, the braking profile may include decreasing the level of braking as vehicle 200 approaches target stopping location 3250. Further, reaching target stopping location 3250 may be accompanied by a release in braking pressure of vehicle 200.

[0340] The foregoing description has been presented for purposes of illustration. It is not exhaustive and is not limited to the precise forms or embodiments disclosed. Modifications and adaptations will be apparent to those skilled in the art from consideration of the specification and practice of the disclosed embodiments. Additionally, although aspects of the disclosed embodiments are described as being stored in memory, one skilled in the art will appreciate that these aspects can also be stored on other types of computer readable media, such as secondary storage devices, for example, hard disks or CD ROM, or other forms of RAM or ROM, USB media, DVD, Blu-ray, or other optical drive media.

[0341] Computer programs based on the written description and disclosed methods are within the skill of an experienced developer. The various programs or program modules can be created using any of the techniques known to one skilled in the art or can be designed in connection with existing software. For example, program sections or program modules can be designed in or by means of .Net Framework, Net Compact Framework (and related languages, such as Visual Basic, C, etc.), Java, C++, Objective-C, HTML, HTML/AJAX combinations, XML, or HTML with included Java applets.

Espacenet

Claims: EP3092599 (A1) — 2016-11-16

SYSTEMS AND METHODS FOR MIMICKING A LEADING VEHICLE

Claims of EP3092599 (B1)

A high quality text as facsimile in your desired language may be available amongst the following family members:

CN106030609 (A) US2015151725 (A1) WO2015083009 (A1) CN110920609 (A) CN111027420 (A) US2015151742 (A1) US2015151751 (A1) US2015151753 (A1) US2015153735 (A1) US2015367850 (A1) US2016052514 (A1) US2016167650 (A1) US2016229410 (A1) US2016347322 (A1) US2019248372 (A1) US2020094839 (A1) US2020108834 (A1) US2020331477 (A1) more

- Original claims
- Claims tree

The EPO does not accept any responsibility for the accuracy of data and information originating from other authorities than the EPO; in particular, the EPO does not guarantee that they are complete, up-to-date or fit for specific purposes.

1. A driver assist navigation system (100) for a primary vehicle (200a), the system comprising:
at least one image capture device (122, 124, 126) configured to acquire a plurality of images of an area in a vicinity of the primary vehicle;
a data interface (128); and
at least one processing device (110) configured to:
receive (1810) the plurality of images via the data interface (128);
determine (1820), from at least some of the plurality of images, a first lane constraint (810) on a first side (802) of the primary vehicle (200a);
determine (1830), from at least some of the plurality of images, a second lane constraint (820) on a second side (812) of the primary vehicle (200a) opposite to the first side (802) of the primary vehicle, wherein the first and second lane constraints define a lane within which the primary vehicle (200a) travels; and
send control signals to the control systems (220, 230, 240) of said primary vehicle for causing (1840) the primary vehicle (200a) to travel within the first and second lane constraints (810, 820);
wherein:
the first lane constraint includes a line marked on a road surface, an edge of a road surface, markers on a road surface, or is estimated based on a determined midpoint of a road surface width;
the second lane constraint includes a line marked on a road surface, an edge of a road surface, markers on a road surface, or is estimated based on a determined midpoint of

a road surface width; and
characterized in that
the at least one processing device (110) is further configured to:
locate (1850) in the plurality of images a leading vehicle (200b);
determine (1860), based on the plurality of images, at least one action taken by the leading vehicle (200b); and
send control signals to the control systems (220, 230, 240) for causing (1870) the primary vehicle (200a) to mimic the at least one action of the leading vehicle (200b);
wherein the at least one action taken by the leading vehicle (200b) includes performing a maneuver such that a first offset distance (c3) on a side of the leading vehicle (200b) adjacent to the first lane constraint (810) is different from a second offset distance (c4) on a side of the leading vehicle (200b) adjacent to the second lane constraint (820), and
wherein mimicking the at least one action includes the at least one processing device causing the primary vehicle to perform a maneuver such that a third offset distance (c1) on a side of the primary vehicle (200a) adjacent to the first lane constraint (810) is different from a fourth offset distance (c2) on a side of the primary vehicle (200a) adjacent to the second lane constraint (820), wherein the third offset distance (c1) equals to the first offset distance (c3) and the fourth offset distance (c2) equals to the second offset distance (c4).

2. The system of claim 1, wherein the at least one action taken by the leading vehicle (200b) further includes a turn, and mimicking the at least one action respectively includes the at least one processing device (110) causing the primary vehicle (200a) to execute a similar turn beginning at a location at which the leading vehicle (200b) initiated its turn.

3. The system of claim 1, wherein the at least one action taken by the leading vehicle further includes an acceleration and wherein mimicking the at least one action respectively includes the at least one processing device causing the primary vehicle to accelerate.

4. The system of claim 1, wherein the at least one action taken by the leading vehicle further includes braking wherein mimicking the at least one action respectively includes the at least one processing device causing the primary vehicle to brake.

5. The system of claim 1, wherein the at least one processing device (110) is further configured to:
determine (1925), based on the plurality of images or based on map data and position information, whether the leading vehicle (200b) is turning at an intersection; and
if the leading vehicle (200b) is determined to be turning at an intersection, decline (1930) to cause the primary vehicle to perform a turn similar to the turn of the leading vehicle (200b).

6. The system of claim 1, wherein the at least one processing device (110) is further configured to:
determine (2025), based on the plurality of images or based on map data and position information, whether the leading vehicle (200b) is turning at an intersection;
determine (2035) whether the leading vehicle (200b) is changing lanes within the intersection;
mimic the turn of the leading vehicle (200b) by causing the primary vehicle (200a) to execute a similar turn beginning at a location at which the leading vehicle (200b)

IPR2025-00943
Tesla EX1002 Page 1667

initiated its turn, if the leading vehicle (200b) is determined to be turning at an intersection and not changing lanes within the intersection; and
decline (2050) to cause the primary vehicle to perform a turn similar to the turn of the leading vehicle (200b) if the leading vehicle (200b) is determined to be turning at an intersection, but changing lanes within the intersection.

7. The system of claim 1, wherein:
the leading vehicle is a first leading vehicle;
at least one processing device (110) is further configured to: locate (2130) a second leading vehicle in the plurality of images; and
cause the primary vehicle (200a) to mimic the at least one action of the first leading vehicle if, based on the plurality of images, the first leading vehicle is determined to be continuing on a path having less of a turning radius than a path followed by the second leading vehicle.

8. The system of claim 1, wherein the at least one processing device (110) is further configured to:
determine (2230) position information for the leading vehicle (200b) and compare (2240) the position information to predetermined map data;
create (2250) a navigation history for the leading vehicle (200b) by tracking movements of the leading vehicle (200b) relative to the predetermined map data; and
determine (2260) whether to cause the primary vehicle (200a) to mimic the at least one action of the leading vehicle (200b) based on the navigation history of the leading vehicle (200b).

9. The system of claim 8, wherein the at least one processing device (110) is configured to:
cause the primary vehicle (200a) to perform a turn similar to a turn of the leading vehicle (200b), even where no turn is detected based on the plurality of images, when the navigation history indicates that the leading vehicle (200b) has successfully navigated at least one prior turn.

10. A primary vehicle for following a leading vehicle, the primary vehicle comprising the system of any preceding claim.

11. A method for navigating a primary vehicle (200a), the method comprising:
acquiring (1810), using at least one image capture device (122, 124, 126), a plurality of images of an area in a vicinity of the primary vehicle (200a);
determining (1820), from at least some of the plurality of images, a first lane constraint (810) on a first side (802) of the primary vehicle (200a);
determining (1830), from at least some of the plurality of images, a second lane constraint (820) on a second side (812) of the primary vehicle (200a) opposite to the first side (802) of the primary vehicle (200a), wherein the first and second lane constraints (810, 820) define a lane within which the primary vehicle (200a) travels; and
sending control signals to the control systems (220, 230, 240) of said primary vehicle for causing (1840) the primary vehicle (200a) to travel within the first and second lane constraints (810, 820); wherein:
the first lane constraint includes a line marked on a road surface, an edge of a road surface, markers on a road surface, or is estimated based on a determined midpoint of a road surface width;
the second lane constraint includes a line marked on a road surface, an edge of a road surface, markers on a road surface, or is estimated based on a determined midpoint of

Tesla EX1002 Page 1668 IPR2025-00943

a road surface width, and
characterized in that the method further comprises:
locating (1850) in the plurality of images a leading vehicle (200b);
determining (1860), based on the plurality of images, at least one action taken by the leading vehicle (200b); and
sending control signals to the control systems (220, 230, 240) for causing (1870) the primary vehicle (200a) to mimic the at least one action of the leading vehicle (200b);
wherein the at least one action taken by the leading vehicle (200b) includes performing a maneuver such that a first offset distance (c3) on a side of the leading vehicle (200b) adjacent to the first lane constraint (810) is different from a second offset distance (c4) on a side of the leading vehicle (200b) adjacent to the second lane constraint (820), and
wherein mimicking the at least one action includes the at least one processing device causing the primary vehicle to perform a maneuver such that a third offset distance (c1) on a side of the primary vehicle (200a) adjacent to the first lane constraint (810) is different from a fourth offset distance (c2) on a side of the primary vehicle (200a) adjacent to the second lane constraint (820), wherein the third offset distance (c1) equals to the first offset distance (c3) and the fourth offset distance (c2) equals to the second offset distance (c4).

12. The method of claim 11, further comprising:
determining (1925), based on the plurality of images or based on map data and position information, whether the leading vehicle (200b) is turning at an intersection; and
if the leading vehicle (200b) is determined to be turning at an intersection, declining (1930) to mimic the turn of the leading vehicle (200b).

13. The method of claim 11, further comprising:
determining (2025), based on the plurality of images or based on map data and position information, whether the leading vehicle (200b) is turning at an intersection;
determining (2035) whether the leading vehicle (200b) is changing lanes within the intersection;
mimicking the turn of the leading vehicle (200b) by causing the primary vehicle (200a) to execute a similar turn beginning at a location at which the leading vehicle (200b) initiated its turn, if the leading vehicle (200b) is determined to be turning at an intersection and not changing lanes within the intersection; and
declining (2050) to cause the primary vehicle to perform a turn similar to the turn of the leading vehicle (200b) if the leading vehicle (200b) is determined to be turning at an intersection, but changing lanes within the intersection.

14. The method of claim 11, wherein the leading vehicle is a first leading vehicle, the method further comprising:
locating (2130) a second leading vehicle in the plurality of images; and
causing the primary vehicle (200a) to mimic the at least one action of first leading vehicle if, based on the plurality of images, the first leading vehicle is determined to be continuing on a path having less of a turning radius than a path followed by the second leading vehicle.

15. The method of claim 11, further comprising to:
determining (2230) position information for the leading vehicle (200b) and compare (2240) the position information to predetermined map data;
creating (2250) a navigation history for the leading vehicle (200b) by tracking

IPR2025-00943
Tesla EX1002 Page 1669

movements of the leading vehicle (200b) relative to the predetermined map data; and determining (2260) whether to cause the primary vehicle (200a) to mimic the at least one action of the leading vehicle (200b) based on the navigation history of the leading vehicle (200b); and

optionally causing the primary vehicle (200a) to mimic a turn of the leading vehicle (200b), even where no turn is detected based on the plurality of images, when the navigation history indicates that the leading vehicle (200b) has successfully navigated at least one prior turn.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Veröffentlichungsnummer:
(11) Publication number: **EP 3 092 599 A0**
(11) Numéro de publication:

Internationale Anmeldung veröffentlicht durch die Weltorganisation für geistiges Eigentum unter der Nummer:

WO 2015/083009 (Art. 153(3) EPÜ).

International application published by the World Intellectual Property Organization under number:

WO 2015/083009 (Art. 153(3) EPC).

Demande internationale publiée par l'Organisation Mondiale de la Propriété Intellectuelle sous le numéro:

WO 2015/083009 (art. 153(3) CBE).

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 3 092 599 B1

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
06.03.2019 Bulletin 2019/10

(21) Application number: **14843184.4**

(22) Date of filing: **04.12.2014**

(51) Int Cl.:
G06K 9/00 *(2006.01)* ***B60K 31/00*** *(2006.01)*

(86) International application number:
PCT/IB2014/003151

(87) International publication number:
WO 2015/083009 (11.06.2015 Gazette 2015/23)

(54) **SYSTEMS AND METHODS FOR MIMICKING A LEADING VEHICLE**

SYSTEME UND VERFAHREN ZUM SIMULIEREN EINES FÜHRUNGSFAHRZEUGS

SYSTEMES ET PROCEDES POUR IMITER UN VEHICULE DE TETE

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR

(30) Priority: **04.12.2013 US 201361911490 P**
14.05.2014 US 201461993084 P
14.05.2014 US 201461993111 P
23.06.2014 US 201462015524 P
09.07.2014 US 201462022221 P
21.08.2014 US 201462040224 P
21.08.2014 US 201462040269 P

(43) Date of publication of application:
16.11.2016 Bulletin 2016/46

(73) Proprietor: **Mobileye Vision Technologies Ltd.**
91450 Jerusalem (IL)

(72) Inventors:
- **CLARKE, Anna**
  9083600 Har Adar (IL)
- **BAGON, Eyal**
  Jerusalem (IL)
- **NIR, Gali**
  Tel Aviv (IL)

(74) Representative: **Finnegan Europe LLP**
1 London Bridge
London SE1 9BG (GB)

(56) References cited:
US-A1- 2011 190 972 US-B1- 8 504 233

- **XIN LIU ET AL: "Vision-based long-distance lane perception and front vehicle location for full autonomous vehicles on highway roads", JOURNAL OF CENTRAL SOUTH UNIVERSITY OF TECHNOLOGY ; SCIENCE & TECHNOLOGY OF MINING AND METALLURGY, CENTRAL SOUTH UNIVERSITY, HEIDELBERG, vol. 19, no. 5, 27 April 2012 (2012-04-27) , pages 1454-1465, XP035048188, ISSN: 1993-0666, DOI: 10.1007/S11771-012-1162-7**
- **NASSER KEHTARNAVAZ ET AL: "A Transportable Neural-Network Approach to Autonomous Vehicle Following", IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, vol. 47, no. 2, 1 May 1998 (1998-05-01), XP011063694, ISSN: 0018-9545**
- **NASSER KEHTARNAVAZ ET AL: "VISUAL CONTROL OF AN AUTONOMOUS VEHICLE (BART)-THE VEHICLE-FOLLOWING PROBLEM", IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, vol. 40, no. 3, 1 August 1991 (1991-08-01) , pages 654-662, XP000266040, ISSN: 0018-9545, DOI: 10.1109/25.97520**

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 3 092 599 B1

Printed by Jouve, 75001 PARIS (FR)

Description

BACKGROUND

I. Technical Field

[0001] The present disclosure relates generally to autonomous vehicle navigation and, more specifically, to systems and methods that use cameras to provide autonomous vehicle navigation features.

II. Background Information

[0002] As technology continues to advance, the goal of a fully autonomous vehicle that is capable of navigating on roadways is on the horizon. Primarily, an autonomous vehicle may be able to identify its environment and navigate without input from a human operator. Autonomous vehicles may also take into account a variety of factors and make appropriate decisions based on those factors to safely and accurately reach an intended destination. For example, various objects-such as other vehicles and pedestrians-are encountered when a vehicle typically travels a roadway. Autonomous driving systems may recognize these objects in a vehicle's environment and take appropriate and timely action to avoid collisions. Additionally, autonomous driving systems may identify other indicators-such as traffic signals, traffic signs, and lane markings-that regulate vehicle movement (e.g., when the vehicle must stop and may go, a speed at which the vehicle must not exceed, where the vehicle must be positioned on the roadway, etc.). Autonomous driving systems may need to determine when a vehicle should change lanes, turn at intersections, change roadways, etc. As is evident from these examples, many factors may need to be addressed in order to provide an autonomous vehicle that is capable of navigating safely and accurately.

[0003] "Vision-based Long-distance Lane Perception and Front Vehicle Location for Full Autonomous Vehicles on Highway Roads" by Xin Liu, et al is directed to a vision-based long-distance lane perception and front location method for decision making by full autonomous vehicles on highway roads.

[0004] "A Transportable Neural-Network Approach to Autonomous Vehicle Following" by Nasser Kehtarnavaz, et al is directed to utilizing neural networks to control the speed and steering of a vehicle for the purpose of performing autonomous vehicle following, while following a leading vehicle.

[0005] US 2011/0190972 A1 is directed to systems for detecting the presence of stationary and non-stationary objects in the vicinity of a travelling vehicle, and controlling vehicle operational parameters in response to presence of such objects.

[0006] US 8,504,233 B1 is directed to a method for controlling a vehicle in an autonomous mode, which includes obtaining lane information that provides an estimated location of a lane on a road on which the vehicle is travelling and, if the lane information has become unavailable or unreliable, using sensor(s) to monitor neighboring vehicle(s) and controlling the vehicle to maintain at least a predetermined minimum distance between the vehicle and the neighboring vehicle(s).

[0007] JP2012-168592 is directed to a method for environment recognition of a vehicle. It comprises recognizing an arrow signal of a traffic light which is located far from the vehicle. It can also comprise performing a preceding vehicle detection and means to detect lanes detection or the road surface.

SUMMARY

[0008] Embodiments consistent with the present disclosure provide systems and methods for autonomous vehicle navigation. The disclosed embodiments may use cameras to provide autonomous vehicle navigation features. For example, consistent with the disclosed embodiments, the disclosed systems may include one, two, or more cameras that monitor the environment of a vehicle and cause a navigational response based on an analysis of images captured by one or more of the cameras.

[0009] In accordance with aspects of the invention there are provided a driver assist navigation system, a vehicle, and a method for navigating a vehicle as set out in claims 1, 10, and 11 respectively.

[0010] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the vehicle; determine from the plurality of images a second lane constraint on a second side of the vehicle opposite to the first side of the vehicle, wherein the first and second lane constraints define a lane within which the vehicle travels and wherein a first distance corresponds to a distance between the first side of the vehicle and the first lane constraint and a second distance corresponds to a distance between the second side of the vehicle and the second lane constraint; determine, based on the plurality of images, whether a lane offset condition exists on the first side of the vehicle; if a lane offset condition exists on the first side of the vehicle, cause the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance; determine, based on the plurality of images, whether a lane offset condition exists on the second side of the vehicle; and if a lane offset condition exists on the second side of the vehicle, cause the vehicle to travel within the first and second lane constraints such that the first distance is less than the second distance.

[0011] Consistent with the disclosed embodiment, the

IPR2025-00943
Tesla EX1002 Page 1673

at least one processing device may be further configured to cause the first distance to be substantially the same as the second distance if, based on the plurality of images, a lane offset condition is determined to exist on both the first and second sides of the vehicle. The at least one processing device may be configured to determine, based on the plurality of images, whether a lane offset condition is substantially non-existent, and if the lane offset condition is substantially non-existent, cause the vehicle to travel within the first and second lane constraints such that the first distance is substantially equal to the second distance. If the driver assist navigation system determines that a lane offset condition exists, the at least one processing device may be configured to cause the first distance to be at least 1.25 times greater than the second distance. Alternatively, the at least one processing device may be further configured to cause the first distance to be at least 2 times greater than the second distance. If the driver assist navigation system determines that the lane offset condition includes a curved road, the at least one processing device may be configured to cause the vehicle to travel within the first and second lane constraints such that the first distance is toward an outside of the curve relative to the vehicle, and the second distance is toward an inside of the curve relative to the vehicle. If the driver assist navigation system determines that the lane offset condition includes a presence of one or more objects on one side of the lane, the at least one processing device may be configured to cause the vehicle to travel within the first and second lane constraints such that the first distance, which is greater than the second distance, is provided toward the one or more objects. The one or more objects may include parked cars, a wall, and/or pedestrians, and may have a height of greater than 10 cm from a road surface. At least one of the first and second lane constraints may include a line marked on a road surface, an edge of a road surface, or a determined midpoint of a road surface width. Determination of whether a lane offset condition exists on the first side of the vehicle may be performed based on a position of a target object on the first side of the vehicle. The causing of the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance may be overridden based on input from a driver of the vehicle.

[0012] Consistent with another disclosed embodiment, a vehicle may include a first vehicle side; a second vehicle side opposite the first vehicle side; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on the first vehicle side; determine from the plurality of images a second lane constraint on the second vehicle side, wherein the first and second lane constraints define a lane within which the vehicle travels and wherein a first distance corresponds to a distance between the first vehicle side and the first lane constraint and a second distance corresponds to a distance between the second vehicle side and the second lane constraint; determine, based on the plurality of images, whether a lane offset condition exists on the first vehicle side; if a lane offset condition exists on the first vehicle side, cause the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance; determine, based on the plurality of images, whether a lane offset condition exists on the second vehicle side; and if a lane offset condition exists on the second vehicle side, cause the vehicle to travel within the first and second lane constraints such that the first distance is less than the second distance.

[0013] Consistent with the disclosed embodiment, the at least one processing device may be further configured to cause the first distance to be substantially the same as the second distance if, based on the plurality of images, a lane offset condition is determined to exist on both the first and second vehicle sides. The at least one processing device may be configured to determine, based on the plurality of images, whether a lane offset condition is substantially non-existent, and if the lane offset condition is substantially non-existent, cause the vehicle to travel within the first and second lane constraints such that the first distance is substantially equal to the second distance. The lane offset condition may include a curved road, and the at least one processing device may be configured to cause the vehicle to travel within the first and second lane constraints such that the first distance is toward an outside of the curve relative to the vehicle, and the second distance is toward an inside of the curve relative to the vehicle. The lane offset condition may include a presence of one or more objects on one side of the lane, and the at least one processing device may be configured to cause the vehicle to travel within the first and second lane constraints such that the first distance, which is greater than the second distance, is provided toward the one or more objects. At least one of the first and second lane constraints may include an edge of a road surface and/or a determined midpoint of a road surface width. Determination of whether a lane offset condition exists on the first side of the vehicle may be performed based on a position of a target object on the first side of the vehicle. The causing of the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance may be overridden based on input from a driver of the vehicle.

[0014] Consistent with another disclosed embodiment, a method is provided for navigating a vehicle. The method may include acquiring, using at least one image capture device, a plurality of images of an area in the vicinity of the vehicle; determining from the plurality of images a first lane constraint on a first side of the vehicle; determining from the plurality of images a second lane constraint on a second side of the vehicle opposite to the

IPR2025-00943
Tesla EX1002 Page 1674

first side of the vehicle, wherein the first and second lane constraints define a lane within which the vehicle travels and wherein a first distance corresponds to a distance between the first side of the vehicle and the first lane constraint and a second distance corresponds to a distance between the second side of the vehicle and the second lane constraint; determining, based on the plurality of images, whether a lane offset condition exists on the first side of the vehicle; if a lane offset condition exists on the first side of the vehicle, causing the vehicle to travel within the first and second lane constraints such that the first distance is greater than the second distance; determining, based on the plurality of images, whether a lane offset condition exists on the second side of the vehicle; and if a lane offset condition exists on the second side of the vehicle, causing the vehicle to travel within the first and second lane constraints such that the first distance is less than the second distance.

[0015] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a current lane of travel from among a plurality of available travel lanes; and cause the vehicle to change lanes if the current lane of travel is not the same as a predetermined default travel lane.

[0016] Consistent with the disclosed embodiment, the predetermined default travel lane may be a right-most lane among the plurality of available travel lanes. The at least one processing device may be configured to activate a turn signal prior to causing the vehicle to change lanes. The at least one processing device may be configured to cause an audible announcement prior to causing the vehicle to change lanes. The at least one processing device may be configured to cause an audible indicator prior to causing the vehicle to change lanes. The predetermined default travel lane may be determined based on the plurality of images of the area in the vicinity of the vehicle, based on data received from a global positioning system, and/or based on an input received from a user.

[0017] Consistent with another disclosed embodiment, a vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a current lane of travel from among a plurality of available travel lanes; and cause the vehicle to change lanes if the current lane of travel is not the same as a predetermined default travel lane.

[0018] Consistent with the disclosed embodiment, the predetermined default travel lane may be a right-most lane among the plurality of available travel lanes. The at least one processing device may be configured to activate a turn signal prior to causing the vehicle to change lanes. The at least one processing device may be configured to cause an audible announcement prior to causing the vehicle to change lanes. The at least one processing device may be configured to cause an audible indicator prior to causing the vehicle to change lanes. The predetermined default travel lane may be determined based on the plurality of images of the area in the vicinity of the vehicle, based on data received from a global positioning system, and/or based on an input received from a user.

[0019] Consistent with another disclosed embodiment, a method is provided for navigating a vehicle. The method may include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the vehicle; determining from the plurality of images a current lane of travel from among a plurality of available travel lanes; and causing the vehicle to change lanes if the current lane of travel is not the same as a predetermined default travel lane.

[0020] Consistent with the disclosed embodiment, the predetermined default travel lane may be a right-most lane among the plurality of available travel lanes. A turn signal may be activated prior to causing the vehicle to change lanes. The predetermined default travel lane may be determined based on the plurality of images of the area in the vicinity of the vehicle.

[0021] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; recognize a curve to be navigated based on map data and vehicle position information; determine an initial target velocity for the vehicle based on at least one characteristic of the curve as reflected in the map data; adjust a velocity of the vehicle to the initial target velocity; determine, based on the plurality of images, one or more observed characteristics of the curve; determine an updated target velocity based on the one or more observed characteristics of the curve; and adjust the velocity of the vehicle to the updated target velocity.

[0022] Consistent with the disclosed embodiment, the at least one characteristic of the curve may include an expected radius of curvature for the curve as determined from the map data. The initial target velocity may be determined based on the expected radius of curvature for the curve and a lateral acceleration limit. The lateral acceleration limit may be between about 0.2 m/sec^2 and 3.0 m/sec^2. Alternatively, the lateral acceleration limit may be between about 1.5 m/sec^2 and 3.0 m/sec^2. The lateral acceleration limit may be adjusted based on a road width determined based on the plurality of images. The one or more observed characteristics of the curve may

IPR2025-00943
Tesla EX1002 Page 1675

include an actual radius of curvature for the curve as determined through analysis of the plurality of images. The updated target velocity may be determined based on the actual radius of curvature for the curve and a lateral acceleration limit. The one or more observed characteristics of the curve may include a rate of change in an actual radius of curvature for the curve as determined through analysis of the plurality of images. The one or more observed characteristics of the curve may include a degree of banking associated with the curve as determined through analysis of the plurality of images. The one or more observed characteristics of the curve may include information, determined through analysis of the plurality of images, relating to the curve conveyed by one or more road signs. The initial target velocity may also be based on a focal range of the at least one image capture device. The initial target velocity may also be based on whether a traffic light is determined to be located along the curve. The initial target velocity may also be based on whether a traffic light is determined to be located within a predetermined distance after the curve.

[0023] Consistent with another disclosed embodiment, a vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; recognize a curve to be navigated based on map data and vehicle position information; determine an initial target velocity for the vehicle based on at least one characteristic of the curve as reflected in the map data; adjust a velocity of the vehicle to the initial target velocity; determine, based on the plurality of images, one or more observed characteristics of the curve; determine an updated target velocity based on the one or more observed characteristics of the curve; and adjust the velocity of the vehicle to the updated target velocity.

[0024] Consistent with the disclosed embodiment, the at least one characteristic of the curve may include an expected radius of curvature for the curve as determined from the map data. The initial target velocity may be determined based on the expected radius of curvature for the curve and a lateral acceleration limit. The lateral acceleration limit may be adjusted based on a road width determined based on the plurality of images. The one or more observed characteristics of the curve may include an actual radius of curvature for the curve as determined through analysis of the plurality of images. The updated target velocity may be determined based on the actual radius of curvature for the curve and a lateral acceleration limit. The one or more observed characteristics of the curve may include a rate of change in an actual radius of curvature for the curve as determined through analysis of the plurality of images. The one or more observed characteristics of the curve may include a degree of banking associated with the curve as determined through analysis of the plurality of images. The one or more observed characteristics of the curve may include information, determined through analysis of the plurality of images, relating to the curve conveyed by one or more road signs. The initial target velocity may also be based on a focal range of the at least one image capture device. The initial target velocity may also be based on whether a traffic light is determined to be located along the curve. The initial target velocity may also be based on whether a traffic light is determined to be located within a predetermined distance after the curve.

[0025] Consistent with another disclosed embodiment, a method is provided for navigating a vehicle. The method may include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the vehicle; recognizing a curve to be navigated based on map data and vehicle position information; determining an initial target velocity for the vehicle based on at least one characteristic of the curve as reflected in the map data; adjusting a velocity of the vehicle to the initial target velocity; determining, based on the plurality of images, one or more observed characteristics of the curve; determining an updated target velocity based on the one or more observed characteristics of the curve; and adjusting the velocity of the vehicle to the updated target velocity.

[0026] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a primary vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the primary vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine, from at least some of the plurality of images, a first lane constraint on a first side of the primary vehicle; determine, from at least some of the plurality of images, a second lane constraint on a second side of the primary vehicle opposite to the first side of the primary vehicle, wherein the first and second lane constraints define a lane within which the primary vehicle travels; cause the primary vehicle to travel within the first and second lane constraints; locate in the plurality of images a leading vehicle; determine, based on the plurality of images, at least one action taken by the leading vehicle; and cause the primary vehicle to mimic the at least one action of the leading vehicle.

[0027] Consistent with the disclosed embodiment, at least one of the first and second lane constraints may include a line marked on a road surface. At least one of the first and second lane constraints may include an edge of a road surface. At least one of the first and second lane constraints may include a determined midpoint of a road surface width. The at least one action taken by the leading vehicle may include a turn, and the at least one processing device may cause the primary vehicle to execute a similar turn beginning at a location at which the leading vehicle initiated its turn. The at least one action taken by the leading vehicle may include an acceleration. The at least one action taken by the leading vehicle may include braking. The at least one action taken by the lead-

IPR2025-00943
Tesla EX1002 Page 1676

ing vehicle may include providing a first offset distance on a side of the leading vehicle adjacent to the first lane constraint that is different from a second offset distance on a side of the leading vehicle adjacent to the second lane constraint. The at least one processing device may be further configured to determine, based on the plurality of images or based on map data and position information, whether the leading vehicle is turning at an intersection, and if the leading vehicle is determined to be turning at an intersection, decline to mimic the turn of the leading vehicle. The at least one processing device may be configured to determine, based on the plurality of images or based on map data and position information, whether the leading vehicle is turning at an intersection, determine whether the leading vehicle is changing lanes within the intersection, mimic the turn of the leading vehicle if the leading vehicle is determined to be turning at an intersection and not changing lanes within the intersection, and decline to mimic the turn of the leading vehicle if the leading vehicle is determined to be turning at an intersection, but changing lanes within the intersection. The at least one processing device may be configured to acquire a first leading vehicle based on the plurality of images, acquire a second leading vehicle based on the plurality of images, and cause the primary vehicle to mimic the first leading vehicle if, based on the plurality of images, the first leading vehicle is determined to be continuing on a path having less of a turning radius than a path followed by the second leading vehicle. The at least one processing device may be configured to determine position information for the leading vehicle and compare the position information to predetermined map data, create a navigation history for the leading vehicle by tracking movements of the leading vehicle relative to the predetermined map data, and determine whether to cause the primary vehicle to mimic the at least one action of the leading vehicle based on the navigation history of the leading vehicle. The at least one processing device may be configured to cause the primary vehicle to mimic a turn of the leading vehicle, even where no turn is detected based on the plurality of images, when the navigation history indicates that the leading vehicle has successfully navigated at least one prior turn.

[0028] Consistent with another disclosed embodiment, a primary vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the primary vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine, from at least some of the plurality of images, a first lane constraint on a first side of the primary vehicle; determine, from at least some of the plurality of images, a second lane constraint on a second side of the primary vehicle opposite to the first side of the primary vehicle, wherein the first and second lane constraints define a lane within which the primary vehicle travels; cause the primary vehicle to travel within the first and second lane constraints; locate in the plurality of images a leading vehicle; determine, based on the plurality of images, at least one action taken by the leading vehicle; and cause the primary vehicle to mimic the at least one action of the leading vehicle.

[0029] Consistent with the disclosed embodiment, at least one of the first and second lane constraints may include a line marked on a road surface. At least one of the first and second lane constraints may include an edge of a road surface. At least one of the first and second lane constraints may include a determined midpoint of a road surface width. The at least one action taken by the leading vehicle may include a turn, and the at least one processing device may cause the primary vehicle to execute a similar turn beginning at a location at which the leading vehicle initiated its turn. The at least one action taken by the leading vehicle may include an acceleration. The at least one action taken by the leading vehicle may include braking. The at least one action taken by the leading vehicle may include providing a first offset distance on a side of the leading vehicle adjacent to the first lane constraint that is different from a second offset distance on a side of the leading vehicle adjacent to the second lane constraint. The at least one processing device may be further configured to determine, based on the plurality of images or based on map data and position information, whether the leading vehicle is turning at an intersection, and if the leading vehicle is determined to be turning at an intersection, decline to mimic the turn of the leading vehicle. The at least one processing device may be configured to determine, based on the plurality of images or based on map data and position information, whether the leading vehicle is turning at an intersection, determine whether the leading vehicle is changing lanes within the intersection, mimic the turn of the leading vehicle if the leading vehicle is determined to be turning at an intersection and not changing lanes within the intersection, and decline to mimic the turn of the leading vehicle if the leading vehicle is determined to be turning at an intersection, but changing lanes within the intersection. The at least one processing device may be configured to acquire a first leading vehicle based on the plurality of images, acquire a second leading vehicle based on the plurality of images, and cause the primary vehicle to mimic the first leading vehicle if, based on the plurality of images, the first leading vehicle is determined to be continuing on a path having less of a turning radius than a path followed by the second leading vehicle. The at least one processing device may be configured to determine position information for the leading vehicle and compare the position information to predetermined map data, create a navigation history for the leading vehicle by tracking movements of the leading vehicle relative to the predetermined map data, and determine whether to cause the primary vehicle to mimic the at least one action of the leading vehicle based on the navigation history of the leading vehicle. The at least one processing device may be configured to cause the primary vehicle to mimic a

IPR2025-00943
Tesla EX1002 Page 1677

turn of the leading vehicle, even where no turn is detected based on the plurality of images, when the navigation history indicates that the leading vehicle has successfully navigated at least one prior turn.

[0030] Consistent with another disclosed embodiment, a method is provided for navigating a primary vehicle. The method ma include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the primary vehicle; determining, from at least some of the plurality of images, a first lane constraint on a first side of the primary vehicle; determining, from at least some of the plurality of images, a second lane constraint on a second side of the primary vehicle opposite to the first side of the primary vehicle, wherein the first and second lane constraints define a lane within which the primary vehicle travels; causing the primary vehicle to travel within the first and second lane constraints; locating in the plurality of images a leading vehicle; determining, based on the plurality of images, at least one action taken by the leading vehicle; and causing the primary vehicle to mimic the at least one action of the leading vehicle.

[0031] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a user vehicle. They system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the user vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the user vehicle; determine from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; acquire, based on the plurality of images, a target vehicle; enable the user vehicle to pass the target vehicle if the target vehicle is determined to be in a lane different from the lane in which the user vehicle is traveling; monitor a position of the target vehicle based on the plurality of images; and cause the user vehicle to abort the pass before completion of the pass, if the target vehicle is determined to be entering the lane in which the user vehicle is traveling.

[0032] Consistent with the disclosed embodiment, causing the user vehicle to abort the pass may include causing an application of brakes in the user vehicle. Causing the user vehicle to abort the pass may include causing an audible announcement. Causing the user vehicle to abort the pass may include causing the user vehicle to stop accelerating. Monitoring the position of the target vehicle may include estimating a speed of the target vehicle. At least one of the first and second lane constraints may include a line marked on a road surface, an edge of a road surface, and/or a determined midpoint of a road surface width.

[0033] Consistent with another disclosed embodiment, a user vehicle may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the user vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the user vehicle; determine from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; acquire, based on the plurality of images, a target vehicle; enable the user vehicle to pass the target vehicle if the target vehicle is determined to be in a lane different from the lane in which the user vehicle is traveling; monitor a position of the target vehicle based on the plurality of images; and cause the user vehicle to abort the pass before completion of the pass, if the target vehicle is determined to be entering the lane in which the user vehicle is traveling.

[0034] Consistent with the disclosed embodiment, causing the user vehicle to abort the pass may include causing an application of brakes in the user vehicle. Causing the user vehicle to abort the pass may include causing an audible announcement. Causing the user vehicle to abort the pass may include causing the user vehicle to stop accelerating. Monitoring the position of the target vehicle may include estimating a speed of the target vehicle. At least one of the first and second lane constraints may include a line marked on a road surface, an edge of a road surface, and/or a determined midpoint of a road surface width.

[0035] Consistent with another disclosed embodiment, a method is provided for navigating a user vehicle. The method may include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the user vehicle; determining from the plurality of images a first lane constraint on a first side of the user vehicle; determining from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; acquiring, based on the plurality of images, a target vehicle; enabling the user vehicle to pass the target vehicle if the target vehicle is determined to be in a lane different from the lane in which the user vehicle is traveling; monitoring a position of the target vehicle based on the plurality of images; and causing the user vehicle to abort the pass before completion of the pass, if the target vehicle is determined to be entering the lane in which the user vehicle is traveling.

[0036] Consistent with the disclosed embodiment, causing the user vehicle to abort the pass may include causing an application of brakes in the user vehicle. Causing the user vehicle to abort the pass may include causing the user vehicle to stop accelerating. At least one of the first and second lane constraints may include a line marked on a road surface, an edge of a road surface, and/or a determined midpoint of a road surface

IPR2025-00943
Tesla EX1002 Page 1678

width.

[0037] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a user vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the user vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the user vehicle; determine from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; determine, based on the plurality of images, whether an encroaching vehicle is approaching from a side of the user vehicle; and cause the user vehicle to maintain a current velocity and to travel within the first and second lane constraints such that a first distance, on the side of the user vehicle that the encroaching vehicle is approaching from, is greater than a second distance.

[0038] Consistent with the disclosed embodiment, the at least one processing device may be further configured to cause the user vehicle to take evasive action if analysis of the plurality of images indicates that the encroaching vehicle has crossed at least one collision threshold. The evasive action may include changing lanes, altering course, and/or braking. The at least one collision threshold may include a minimum predetermined distance to maintain between the user vehicle and the encroaching vehicle. The at least one collision threshold may include a time to collision. The at least one processing device may be configured to cause the user vehicle to return to a center of the lane, within which the user vehicle originally traveled, after completion of the evasive action. The at least one image capture device may have a field of view of at least 140 degrees. The at least one processing device may be configured to cause the user vehicle to resume a preset speed after braking in response to an encroaching vehicle. The at least one processing device may be configured to cause the user vehicle to return to a lane location within the first and second lane constraints such that the first distance between the user vehicle and the first lane constraint is substantially equal to the second distance between the user vehicle and the second lane constraint after altering course in response to an encroaching vehicle.

[0039] Consistent with another disclosed embodiment, a user vehicle may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the user vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; determine from the plurality of images a first lane constraint on a first side of the user vehicle; determine from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; determine, based on the plurality of images, whether an encroaching vehicle is approaching from a side of the user vehicle; and cause the user vehicle to maintain a current velocity and to travel within the first and second lane constraints such that a first distance, on the side of the user vehicle that the encroaching vehicle is approaching from, is greater than a second distance.

[0040] Consistent with the disclosed embodiment, the at least one processing device may be further configured to cause the user vehicle to take evasive action if analysis of the plurality of images indicates that the encroaching vehicle has crossed at least one collision threshold. The evasive action may include changing lanes, altering course, and/or braking. The at least one collision threshold may include a minimum predetermined distance to maintain between the user vehicle and the encroaching vehicle. The at least one collision threshold may include a time to collision. The at least one processing device may be configured to cause the user vehicle to return to a center of the lane, within which the user vehicle originally traveled, after completion of the evasive action. The at least one image capture device may have a field of view of at least 140 degrees. The at least one processing device may be configured to cause the user vehicle to resume a preset speed after braking in response to an encroaching vehicle. The at least one processing device may be configured to cause the user vehicle to return to a lane location within the first and second lane constraints such that the first distance between the user vehicle and the first lane constraint is substantially equal to the second distance between the user vehicle and the second lane constraint after altering course in response to an encroaching vehicle.

[0041] Consistent with another disclosed embodiment, a method for navigating a user vehicle is provided. The method may include acquiring, using at least one image capture device, a plurality of images of an area in a vicinity of the user vehicle; determining from the plurality of images a first lane constraint on a first side of the user vehicle; determining from the plurality of images a second lane constraint on a second side of the user vehicle opposite to the first side of the user vehicle, wherein the first and second lane constraints define a lane within which the user vehicle travels; determining, based on the plurality of images, whether an encroaching vehicle is approaching from a side of the user vehicle; and causing the user vehicle to maintain a current velocity and to travel within the first and second lane constraints such that a first distance, on the side of the user vehicle that the encroaching vehicle is approaching from, is greater than a second distance.

[0042] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a primary vehicle. They system may include at least one image capture device configured to acquire a plurality of images of

IPR2025-00943
Tesla EX1002 Page 1679

an area in a vicinity of the primary vehicle; a data interface; and at least one processing device. The at least processing device may be configured to: receive the plurality of images via the data interface; identify a target object within the plurality of images; monitor, via the plurality of images, a motion of the target object and a distance between the primary vehicle and the target object; determine an indicator of an intercept time between the primary vehicle and the target object based on the monitored motion and the distance between the primary vehicle and the target object; and cause a response in the primary vehicle based on a comparison of the intercept time to a plurality of predetermined intercept thresholds.

[0043] Consistent with the disclosed embodiment, the target object may be a target vehicle traveling in a lane different from the primary vehicle. The intercept time may be based on an indicator of speed of the tracked object, as determined based on the plurality of images. The intercept time may be based on an indicator of an acceleration of the tracked object, as determined based on the plurality of images. The at least one processing device may be further configured to determine whether the target object is traveling in the same lane as the primary vehicle, and forego the response in the primary vehicle if the target object is determined to be traveling in a lane different from the primary vehicle. Determination of whether the target object is traveling in the same lane as the primary vehicle may be based on the plurality of images. Determination of whether the target object is traveling in the same lane as the primary vehicle may be based on at least one of predetermined map data or yaw rate of the primary vehicle. The response may be based on an averaged motion profile associated with the target object if the indicator of intercept time indicates that the target object is outside a first intercept threshold relative to the vehicle. The response may be based on actual motion of the target object if the indicator of intercept time indicates that the target object is between the first intercept threshold and a second intercept threshold relative to the vehicle. The response may include an emergency avoidance action if the indicator of intercept time indicates that the target object is within the second intercept threshold. The emergency avoidance action may include changing lanes and/or emergency braking.

[0044] Consistent with another disclosed embodiment, a primary vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the primary vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; identify a target object within the plurality of images; monitor, via the plurality of images, a motion of the target object and a distance between the primary vehicle and the target object; determine an indicator of an intercept time between the primary vehicle and the target object based on the monitored motion and the distance between the primary vehicle and the target object; and cause a response in the primary vehicle based on a comparison of the intercept time to a plurality of predetermined intercept thresholds.

[0045] Consistent with the disclosed embodiment, the target object may be a target vehicle traveling in a lane different from the primary vehicle. The intercept time may be based on an indicator of speed of the tracked object, as determined based on the plurality of images. The intercept time may be based on an indicator of an acceleration of the tracked object, as determined based on the plurality of images. The at least one processing device may be further configured to determine whether the target object is traveling in the same lane as the primary vehicle, and forego the response in the primary vehicle if the target object is determined to be traveling in a lane different from the primary vehicle. Determination of whether the target object is traveling in the same lane as the primary vehicle may be based on the plurality of images. Determination of whether the target object is traveling in the same lane as the primary vehicle may be based on at least one of predetermined map data or yaw rate of the primary vehicle. The response may be based on an averaged motion profile associated with the target object if the indicator of intercept time indicates that the target object is outside a first intercept threshold relative to the vehicle. The response may be based on actual motion of the target object if the indicator of intercept time indicates that the target object is between the first intercept threshold and a second intercept threshold relative to the vehicle. The response may include an emergency avoidance action if the indicator of intercept time indicates that the target object is within the second intercept threshold. The emergency avoidance action may include changing lanes and/or emergency braking.

[0046] Consistent with another disclosed embodiment, a method for navigating a primary vehicle may include acquiring, via at least one image capture device, a plurality of images of an area in a vicinity of the primary vehicle; identifying a target object within the plurality of images; monitoring, based on the plurality of images, a motion of the target object and a distance between the primary vehicle and the target object; determining an indicator of an intercept time between the primary vehicle and the target object based on the monitored motion and the distance between the primary vehicle and the target object; and causing a response in the primary vehicle based on a comparison of the intercept time to a plurality of predetermined intercept thresholds.

[0047] Consistent with a disclosed embodiment, a driver assist navigation system is provided for a vehicle. The system may include at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; identify, based on analysis of the plurality of images, a trigger for stopping the vehicle; and based on the identified trigger, cause the vehicle to stop according to a braking profile including a first segment

IPR2025-00943
Tesla EX1002 Page 1680

associated with a first deceleration rate, a second segment which includes a second deceleration rate less than the first deceleration rate, and a third segment in which a level of braking is decreased as a target stopping location is approached, as determined based on the analysis of the plurality of images.

[0048] Consistent with the disclosed embodiment, reaching of the target stopping location may be accompanied by a release in braking pressure. The trigger may include an object. The object may include another vehicle in front of the vehicle. The trigger may include a traffic light signal, a stop sign, and/or a speed limit sign. The at least one processing device may be further configured to provide an audible announcement identifying the trigger. The at least one processing device may be configured to cause an audible announcement indicating that the vehicle is braking.

[0049] Consistent with another disclosed embodiment, a vehicle may include a body; at least one image capture device configured to acquire a plurality of images of an area in a vicinity of the vehicle; a data interface; and at least one processing device. The at least one processing device may be configured to: receive the plurality of images via the data interface; identify, based on analysis of the plurality of images, a trigger for stopping the vehicle; and based on the identified trigger, cause the vehicle to stop according to a braking profile including a first segment associated with a first deceleration rate, a second segment which includes a second deceleration rate less than the first deceleration rate, and a third segment in which a level of braking is decreased as a target stopping location is approached, as determined based on the analysis of the plurality of images.

[0050] Consistent with the disclosed embodiment, reaching of the target stopping location may be accompanied by a release in braking pressure. The trigger may include an object. The object may include another vehicle in front of the vehicle. The trigger may include a traffic light signal, a stop sign, and/or a speed limit sign. The at least one processing device may be further configured to provide an audible announcement identifying the trigger. The at least one processing device may be configured to cause an audible announcement indicating that the vehicle is braking.

[0051] Consistent with another disclosed embodiment, a method is provided for navigating a vehicle. The method may include acquiring, via at least one image capture plurality of images of an area in a vicinity of the vehicle; identifying, based on analysis of the plurality of images, a trigger for stopping the vehicle; and based on the identified trigger, causing the vehicle to stop according to a braking profile including a first segment associated with a first deceleration rate, a second segment which includes a second deceleration rate less than the first deceleration rate, and a third segment in which a level of braking is decreased as a target stopping location is approached, as determined based on the analysis of the plurality of images.

[0052] Consistent with other disclosed embodiments, non-transitory computer-readable storage media may store program instructions, which are executed by at least one processing device and perform any of the methods described herein.

[0053] The foregoing general description and the following detailed description are exemplary and explanatory.

BRIEF DESCRIPTION OF THE DRAWINGS

[0054] The accompanying drawings, which are incorporated in and constitute a part of this disclosure, illustrate various disclosed embodiments. In the drawings:

FIG. 1 is a diagrammatic representation of an exemplary system consistent with the disclosed embodiments.

FIG. 2A is a diagrammatic side view representation of an exemplary vehicle including a system consistent with the disclosed embodiments.

FIG. 2B is a diagrammatic top view representation of the vehicle and system shown in FIG. 2A consistent with the disclosed embodiments.

FIG. 2C is a diagrammatic top view representation of another embodiment of a vehicle including a system consistent with the disclosed embodiments.

FIG. 2D is a diagrammatic top view representation of yet another embodiment of a vehicle including a system consistent with the disclosed embodiments.

FIG. 2E is a diagrammatic top view representation of yet another embodiment of a vehicle including a system consistent with the disclosed embodiments.

FIG. 2F is a diagrammatic representation of exemplary vehicle control systems consistent with the disclosed embodiments.

FIG. 3A is a diagrammatic representation of an interior of a vehicle including a rearview mirror and a user interface for a vehicle imaging system consistent with the disclosed embodiments.

FIG. 3B is an illustration of an example of a camera mount that is configured to be positioned behind a rearview mirror and against a vehicle windshield consistent with the disclosed embodiments.

FIG. 3C is an illustration of the camera mount shown in FIG. 3B from a different perspective consistent with the disclosed embodiments.

FIG. 3D is an illustration of an example of a camera mount that is configured to be positioned behind a rearview mirror and against a vehicle windshield consistent with the disclosed embodiments.

FIG. 4 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

FIG. 5A is a flowchart showing an exemplary process for causing one or more navigational responses based on monocular image analysis consistent with

IPR2025-00943
Tesla EX1002 Page 1681

disclosed embodiments.

FIG. 5B is a flowchart showing an exemplary process for detecting one or more vehicles and/or pedestrians in a set of images consistent with the disclosed embodiments.

FIG. 5C is a flowchart showing an exemplary process for detecting road marks and/or lane geometry information in a set of images consistent with the disclosed embodiments.

FIG. 5D is a flowchart showing an exemplary process for detecting traffic lights in a set of images consistent with the disclosed embodiments.

FIG. 5E is a flowchart showing an exemplary process for causing one or more navigational responses based on a vehicle path consistent with the disclosed embodiments.

FIG. 5F is a flowchart showing an exemplary process for determining whether a leading vehicle is changing lanes consistent with the disclosed embodiments.

FIG. 6 is a flowchart showing an exemplary process for causing one or more navigational responses based on stereo image analysis consistent with the disclosed embodiments.

FIG. 7 is a flowchart showing an exemplary process for causing one or more navigational responses based on an analysis of three sets of images consistent with the disclosed embodiments.

FIG. 8A is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint in which a lane offset condition exists on the first vehicle side consistent with the disclosed embodiments.

FIG. 8B is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint in which a lane offset condition exists on the second vehicle side consistent with the disclosed embodiments.

FIG. 8C is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint on a curved road consistent with the disclosed embodiments.

FIG. 8D is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint in which a lane offset condition exists on both the first vehicle side and on the second vehicle side consistent with the disclosed embodiments.

FIG. 9 is a diagrammatic representation of the memory of an exemplary navigation system consistent with the disclosed embodiments.

FIG. 10 is a flowchart of an exemplary method for navigating a vehicle consistent with the disclosed embodiments.

FIG 11 is a diagrammatic representation of an exemplary vehicle including a navigation system traveling within a default lane consistent with the disclosed embodiments.

FIG 12A is a diagrammatic representation of the memory of an exemplary navigation system consistent with the disclosed embodiments.

FIG 12B is a flowchart of an exemplary process for navigating a vehicle consistent with the disclosed embodiments.

FIGS. 13A and 13B are diagrammatic representations of an exemplary vehicle approaching, and navigating, a curve with one or more characteristics consistent with the disclosed embodiments.

FIG. 14 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

FIG. 15 is a flow chart showing an exemplary process for controlling the velocity of a vehicle based on a detected curve and observed characteristics of the curve consistent with disclosed embodiments.

FIGS. 16A-16D are diagrammatic representations of a primary vehicle mimicking one or more actions of a leading vehicle consistent with the disclosed embodiments.

FIG. 17 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

FIG. 18 is a flow chart showing an exemplary process for causing a primary vehicle to mimic one or more actions of a leading vehicle consistent with disclosed embodiments.

FIG. 19 is a flow chart showing an exemplary process for causing a primary vehicle to decline to mimic a turn of a leading vehicle consistent with disclosed embodiments.

FIG. 20 is a flow chart showing an exemplary process for causing a primary vehicle to mimic or decline a turn of a leading vehicle consistent with disclosed embodiments.

FIG. 21 is a flow chart showing another exemplary process for causing a primary vehicle to mimic one or more actions of a leading vehicle consistent with disclosed embodiments.

FIG. 22 is a flow chart showing an exemplary process for causing a primary vehicle to mimic one or more actions of a leading vehicle based on a navigation history consistent with disclosed embodiments.

FIG. 23 is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint on a roadway along on which another vehicle is traveling consistent with the disclosed embodiments.

FIG. 24A is a diagrammatic representation of a memory storing instructions for navigating a vehicle among encroaching vehicles consistent with the disclosed embodiments.

FIG. 24B is a flowchart of an exemplary process for navigating a vehicle among encroaching vehicles consistent with the disclosed embodiments.

IPR2025-00943
Tesla EX1002 Page 1682

FIG. 25 is a diagrammatic representation of an exemplary vehicle traveling within a first and a second lane constraint on a roadway along on which another vehicle is traveling consistent with the disclosed embodiments.

FIG. 26 is a diagrammatic representation of a memory storing instructions for detecting and responding to traffic laterally encroaching on a vehicle consistent with the disclosed embodiments.

FIG. 27 is a flowchart of an exemplary process for detecting and responding to traffic laterally encroaching on a vehicle consistent with the disclosed embodiments.

FIGS. 28A and 28B are diagrammatic representations of causing a response in a primary vehicle consistent with the disclosed embodiments.

FIG. 29 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments.

FIG. 30 is a flow chart showing an exemplary process for causing a response in a primary vehicle consistent with disclosed embodiments.

FIG. 31 is a flow chart showing an exemplary process for declining to cause a response in a primary vehicle consistent with disclosed embodiments.

FIG. 32A is a diagrammatic representation of an exemplary vehicle in which an object is present in the vicinity of vehicle consistent with the disclosed embodiments.

FIG. 32B illustrates an exemplary braking profile, consistent with the disclosed embodiments.

FIG. 33 is a diagrammatic representation of the memory of an exemplary navigation system consistent with the disclosed embodiments.

FIG. 34 is a flowchart of an exemplary method for implementing a braking profile while navigating a vehicle consistent with the disclosed embodiments.

DETAILED DESCRIPTION

[0055] The following detailed description refers to the accompanying drawings. Wherever possible, the same reference numbers are used in the drawings and the following description to refer to the same or similar parts. While several illustrative embodiments are described herein, modifications, adaptations and other implementations are possible. For example, substitutions, additions or modifications may be made to the components illustrated in the drawings, and the illustrative methods described herein may be modified by substituting, reordering, removing, or adding steps to the disclosed methods. Accordingly, the following detailed description is not limited to the disclosed embodiments and examples. Instead, the proper scope is defined by the appended claims.

[0056] Disclosed embodiments provide systems and methods that use cameras to provide autonomous navigation features. In various embodiments, the system may include one, two or more cameras that monitor the environment of a vehicle. In one embodiment, the system may monitor and adjust the free space between a vehicle and the boundaries of the lane within which the vehicle is traveling. In another embodiment, the system may select a particular lane as a default lane for the vehicle to use while traveling. In another embodiment, the system may control the speed of the vehicle in different scenarios, such as while making a turn. In yet another embodiment, the system may mimic the actions of a leading vehicle. In yet another embodiment, the system may monitor a target vehicle and enable the vehicle to pass the target vehicle under certain conditions (e.g., if the target vehicle is traveling in a lane different from the lane within which the vehicle is traveling. In still yet another embodiment, the system may produce a natural feeling response to a laterally encroaching vehicle, such as a vehicle attempting to move into the lane within which the vehicle is traveling.

[0057] FIG. 1 is a block diagram representation of a system 100 consistent with the exemplary disclosed embodiments. System 100 may include various components depending on the requirements of a particular implementation. In some embodiments, system 100 may include a processing unit 110, an image acquisition unit 120, a position sensor 130, one or more memory units 140, 150, a map database 160, and a user interface 170. Processing unit 110 may include one or more processing devices. In some embodiments, processing unit 110 may include an applications processor 180, an image processor 190, or any other suitable processing device. Similarly, image acquisition unit 120 may include any number of image acquisition devices and components depending on the requirements of a particular application. In some embodiments, image acquisition unit 120 may include one or more image capture devices (e.g., cameras), such as image capture device 122, image capture device 124, and image capture device 126. System 100 may also include a data interface 128 communicatively connecting processing device 110 to image acquisition device 120. For example, data interface 128 may include any wired and/or wireless link or links for transmitting image data acquired by image accusation device 120 to processing unit 110.

[0058] Both applications processor 180 and image processor 190 may include various types of processing devices. For example, either or both of applications processor 180 and image processor 190 may include a microprocessor, preprocessors (such as an image preprocessor), graphics processors, a central processing unit (CPU), support circuits, digital signal processors, integrated circuits, memory, or any other types of devices suitable for running applications and for image processing and analysis. In some embodiments, applications processor 180 and/or image processor 190 may include any type of single or multi-core processor, mobile device microcontroller, central processing unit, etc. Various

IPR2025-00943
Tesla EX1002 Page 1683

processing devices may be used, including, for example, processors available from manufacturers such as Intel®, AMD®, etc. and may include various architectures (e.g., x86 processor, ARM®, etc.).

[0059] In some embodiments, applications processor 180 and/or image processor 190 may include any of the EyeQ series of processor chips available from Mobileye®. These processor designs each include multiple processing units with local memory and instruction sets. Such processors may include video inputs for receiving image data from multiple image sensors and may also include video out capabilities. In one example, the EyeQ2® uses 90nm-micron technology operating at 332Mhz. The EyeQ2® architecture consists of two floating point, hyper-thread 32-bit RISC CPUs (MIPS32® 34K® cores), five Vision Computing Engines (VCE), three Vector Microcode Processors (VMP®), Denali 64-bit Mobile DDR Controller, 128-bit internal Sonics Interconnect, dual 16-bit Video input and 18-bit Video output controllers, 16 channels DMA and several peripherals. The MIPS34K CPU manages the five VCEs, three VMP™ and the DMA, the second MIPS34K CPU and the multi-channel DMA as well as the other peripherals. The five VCEs, three VMP® and the MIPS34K CPU can perform intensive vision computations required by multi-function bundle applications. In another example, the EyeQ3®, which is a third generation processor and is six times more powerful that the EyeQ2®, may be used in the disclosed embodiments.

[0060] While FIG. 1 depicts two separate processing devices included in processing unit 110, more or fewer processing devices may be used. For example, in some embodiments, a single processing device may be used to accomplish the tasks of applications processor 180 and image processor 190. In other embodiments, these tasks may be performed by more than two processing devices.

[0061] Processing unit 110 may comprise various types of devices. For example, processing unit 110 may include various devices, such as a controller, an image preprocessor, a central processing unit (CPU), support circuits, digital signal processors, integrated circuits, memory, or any other types of devices for image processing and analysis. The image preprocessor may include a video processor for capturing, digitizing and processing the imagery from the image sensors. The CPU may comprise any number of microcontrollers or microprocessors. The support circuits may be any number of circuits generally well known in the art, including cache, power supply, clock and input-output circuits. The memory may store software that, when executed by the processor, controls the operation of the system. The memory may include databases and image processing software. The memory may comprise any number of random access memories, read only memories, flash memories, disk drives, optical storage, tape storage, removable storage and other types of storage. In one instance, the memory may be separate from the processing unit 110. In another instance, the memory may be integrated into the processing unit 110.

[0062] Each memory 140, 150 may include software instructions that when executed by a processor (e.g., applications processor 180 and/or image processor 190), may control operation of various aspects of system 100. These memory units may include various databases and image processing software. The memory units may include random access memory, read only memory, flash memory, disk drives, optical storage, tape storage, removable storage and/or any other types of storage. In some embodiments, memory units 140, 150 may be separate from the applications processor 180 and/or image processor 190. In other embodiments, these memory units may be integrated into applications processor 180 and/or image processor 190.

[0063] Position sensor 130 may include any type of device suitable for determining a location associated with at least one component of system 100. In some embodiments, position sensor 130 may include a GPS receiver. Such receivers can determine a user position and velocity by processing signals broadcasted by global positioning system satellites. Position information from position sensor 130 may be made available to applications processor 180 and/or image processor 190.

[0064] User interface 170 may include any device suitable for providing information to or for receiving inputs from one or more users of system 100. In some embodiments, user interface 170 may include user input devices, including, for example, a touchscreen, microphone, keyboard, pointer devices, track wheels, cameras, knobs, buttons, etc. With such input devices, a user may be able to provide information inputs or commands to system 100 by typing instructions or information, providing voice commands, selecting menu options on a screen using buttons, pointers, or eye-tracking capabilities, or through any other suitable techniques for communicating information to system 100.

[0065] User interface 170 may be equipped with one or more processing devices configured to provide and receive information to or from a user and process that information for use by, for example, applications processor 180. In some embodiments, such processing devices may execute instructions for recognizing and tracking eye movements, receiving and interpreting voice commands, recognizing and interpreting touches and/or gestures made on a touchscreen, responding to keyboard entries or menu selections, etc. In some embodiments, user interface 170 may include a display, speaker, tactile device, and/or any other devices for providing output information to a user.

[0066] Map database 160 may include any type of database for storing map data useful to system 100. In some embodiments, map database 160 may include data relating to the position, in a reference coordinate system, of various items, including roads, water features, geographic features, businesses, points of interest, restaurants, gas stations, etc. Map database 160 may store not

IPR2025-00943
Tesla EX1002 Page 1684

only the locations of such items, but also descriptors relating to those items, including, for example, names associated with any of the stored features. In some embodiments, map database 160 may be physically located with other components of system 100. Alternatively or additionally, map database 160 or a portion thereof may be located remotely with respect to other components of system 100 (e.g., processing unit 110). In such embodiments, information from map database 160 may be downloaded over a wired or wireless data connection to a network (e.g., over a cellular network and/or the Internet, etc.).

[0067] Image capture devices 122, 124, and 126 may each include any type of device suitable for capturing at least one image from an environment. Moreover, any number of image capture devices may be used to acquire images for input to the image processor. Some embodiments may include only a single image capture device, while other embodiments may include two, three, or even four or more image capture devices. Image capture devices 122,124, and 126 will be further described with reference to FIGS. 2B-2E, below.

[0068] System 100, or various components thereof, may be incorporated into various different platforms. In some embodiments, system 100 may be included on a vehicle 200, as shown in FIG. 2A. For example, vehicle 200 may be equipped with a processing unit 110 and any of the other components of system 100, as described above relative to FIG. 1. While in some embodiments vehicle 200 may be equipped with only a single image capture device (e.g., camera), in other embodiments, such as those discussed in connection with FIGS. 2B-2E, multiple image capture devices may be used. For example, either of image capture devices 122 and 124 of vehicle 200, as shown in FIG. 2A, may be part of an ADAS (Advanced Driver Assistance Systems) imaging set.

[0069] The image capture devices included on vehicle 200 as part of the image acquisition unit 120 may be positioned at any suitable location. In some embodiments, as shown in FIGS. 2A-2E and 3A-3C, image capture device 122 may be located in the vicinity of the rearview mirror. This position may provide a line of sight similar to that of the driver of vehicle 200, which may aid in determining what is and is not visible to the driver. Image capture device 122 may be positioned at any location near the rearview mirror, but placing image capture device 122 on the driver side of the mirror may further aid in obtaining images representative of the driver's field of view and/or line of sight.

[0070] Other locations for the image capture devices of image acquisition unit 120 may also be used. For example, image capture device 124 may be located on or in a bumper of vehicle 200. Such a location may be especially suitable for image capture devices having a wide field of view. The line of sight of bumper-located image capture devices can be different from that of the driver and, therefore, the bumper image capture device and driver may not always see the same objects. The image capture devices (e.g., image capture devices 122, 124, and 126) may also be located in other locations. For example, the image capture devices may be located on or in one or both of the side mirrors of vehicle 200, on the roof of vehicle 200, on the hood of vehicle 200, on the trunk of vehicle 200, on the sides of vehicle 200, mounted on, positioned behind, or positioned in front of any of the windows of vehicle 200, and mounted in or near light figures on the front and/or back of vehicle 200, etc.

[0071] In addition to image capture devices, vehicle 200 may include various other components of system 100. For example, processing unit 110 may be included on vehicle 200 either integrated with or separate from an engine control unit (ECU) of the vehicle. Vehicle 200 may also be equipped with a position sensor 130, such as a GPS receiver and may also include a map database 160 and memory units 140 and 150.

[0072] FIG. 2A is a diagrammatic side view representation of an exemplary vehicle imaging system consistent with the disclosed embodiments. FIG. 2B is a diagrammatic top view illustration of the embodiment shown in FIG. 2A. As illustrated in FIG. 2B, the disclosed embodiments may include a vehicle 200 including in its body a system 100 with a first image capture device 122 positioned in the vicinity of the rearview mirror and/or near the driver of vehicle 200, a second image capture device 124 positioned on or in a bumper region (e.g., one of bumper regions 210) of vehicle 200, and a processing unit 110.

[0073] As illustrated in FIG. 2C, image capture devices 122 and 124 may both be positioned in the vicinity of the rearview mirror and/or near the driver of vehicle 200. Additionally, while two image capture devices 122 and 124 are shown in FIGS. 2B and 2C, it should be understood that other embodiments may include more than two image capture devices. For example, in the embodiments shown in FIGS. 2D and 2E, first, second, and third image capture devices 122, 124, and 126, are included in the system 100 of vehicle 200.

[0074] As illustrated in FIG. 2D, image capture device 122 may be positioned in the vicinity of the rearview mirror and/or near the driver of vehicle 200, and image capture devices 124 and 126 may be positioned on or in a bumper region (e.g., one of bumper regions 210) of vehicle 200. And as shown in FIG. 2E, image capture devices 122, 124, and 126 may be positioned in the vicinity of the rearview mirror and/or near the driver seat of vehicle 200. The disclosed embodiments are not limited to any particular number and configuration of the image capture devices, and the image capture devices may be positioned in any appropriate location within and/or on vehicle 200.

[0075] It is to be understood that the disclosed embodiments are not limited to vehicles and could be applied in other contexts. It is also to be understood that disclosed embodiments are not limited to a particular type of vehicle 200 and may be applicable to all types of vehicles includ-

IPR2025-00943
Tesla EX1002 Page 1685

ing automobiles, trucks, trailers, and other types of vehicles.

[0076] The first image capture device 122 may include any suitable type of image capture device. Image capture device 122 may include an optical axis. In one instance, the image capture device 122 may include an Aptina M9V024 WVGA sensor with a global shutter. In other embodiments, image capture device 122 may provide a resolution of 1280x960 pixels and may include a rolling shutter. Image capture device 122 may include various optical elements. In some embodiments one or more lenses may be included, for example, to provide a desired focal length and field of view for the image capture device. In some embodiments, image capture device 122 may be associated with a 6mm lens or a 12mm lens. In some embodiments, image capture device 122 may be configured to capture images having a desired field-of-view (FOV) 202, as illustrated in FIG. 2D. For example, image capture device 122 may be configured to have a regular FOV, such as within a range of 40 degrees to 56 degrees, including a 46 degree FOV, 50 degree FOV, 52 degree FOV, or greater. Alternatively, image capture device 122 may be configured to have a narrow FOV in the range of 23 to 40 degrees, such as a 28 degree FOV or 36 degree FOV. In addition, image capture device 122 may be configured to have a wide FOV in the range of 100 to 180 degrees. In some embodiments, image capture device 122 may include a wide angle bumper camera or one with up to a 180 degree FOV.

[0077] The first image capture device 122 may acquire a plurality of first images relative to a scene associated with the vehicle 200. Each of the plurality of first images may be acquired as a series of image scan lines, which may be captured using a rolling shutter. Each scan line may include a plurality of pixels.

[0078] The first image capture device 122 may have a scan rate associated with acquisition of each of the first series of image scan lines. The scan rate may refer to a rate at which an image sensor can acquire image data associated with each pixel included in a particular scan line.

[0079] Image capture devices 122, 124, and 126 may contain any suitable type and number of image sensors, including CCD sensors or CMOS sensors, for example. In one embodiment, a CMOS image sensor may be employed along with a rolling shutter, such that each pixel in a row is read one at a time, and scanning of the rows proceeds on a row-by-row basis until an entire image frame has been captured. In some embodiments, the rows may be captured sequentially from top to bottom relative to the frame.

[0080] The use of a rolling shutter may result in pixels in different rows being exposed and captured at different times, which may cause skew and other image artifacts in the captured image frame. On the other hand, when the image capture device 122 is configured to operate with a global or synchronous shutter, all of the pixels may be exposed for the same amount of time and during a common exposure period. As a result, the image data in a frame collected from a system employing a global shutter represents a snapshot of the entire FOV (such as FOV 202) at a particular time. In contrast, in a rolling shutter application, each row in a frame is exposed and data is capture at different times. Thus, moving objects may appear distorted in an image capture device having a rolling shutter. This phenomenon will be described in greater detail below.

[0081] The second image capture device 124 and the third image capturing device 126 may be any type of image capture device. Like the first image capture device 122, each of image capture devices 124 and 126 may include an optical axis. In one embodiment, each of image capture devices 124 and 126 may include an Aptina M9V024 WVGA sensor with a global shutter. Alternatively, each of image capture devices 124 and 126 may include a rolling shutter. Like image capture device 122, image capture devices 124 and 126 may be configured to include various lenses and optical elements. In some embodiments, lenses associated with image capture devices 124 and 126 may provide FOVs (such as FOVs 204 and 206) that are the same as, or narrower than, a FOV (such as FOV 202) associated with image capture device 122. For example, image capture devices 124 and 126 may have FOVs of 40 degrees, 30 degrees, 26 degrees, 23 degrees, 20 degrees, or less.

[0082] Image capture devices 124 and 126 may acquire a plurality of second and third images relative to a scene associated with the vehicle 200. Each of the plurality of second and third images may be acquired as a second and third series of image scan lines, which may be captured using a rolling shutter. Each scan line or row may have a plurality of pixels. Image capture devices 124 and 126 may have second and third scan rates associated with acquisition of each of image scan lines included in the second and third series.

[0083] Each image capture device 122, 124, and 126 may be positioned at any suitable position and orientation relative to vehicle 200. The relative positioning of the image capture devices 122, 124, and 126 may be selected to aid in fusing together the information acquired from the image capture devices. For example, in some embodiments, a FOV (such as FOV 204) associated with image capture device 124 may overlap partially or fully with a FOV (such as FOV 202) associated with image capture device 122 and a FOV (such as FOV 206) associated with image capture device 126.

[0084] Image capture devices 122, 124, and 126 may be located on vehicle 200 at any suitable relative heights. In one instance, there may be a height difference between the image capture devices 122, 124, and 126, which may provide sufficient parallax information to enable stereo analysis. For example, as shown in FIG. 2A, the two image capture devices 122 and 124 are at different heights. There may also be a lateral displacement difference between image capture devices 122, 124, and 126, giving additional parallax information for stereo anal-

IPR2025-00943
Tesla EX1002 Page 1686

ysis by processing unit 110, for example. The difference in the lateral displacement may be denoted by d_x, as shown in FIGS. 2C and 2D. In some embodiments, fore or aft displacement (e.g., range displacement) may exist between image capture devices 122, 124, and 126. For example, image capture device 122 may be located 0.5 to 2 meters or more behind image capture device 124 and/or image capture device 126. This type of displacement may enable one of the image capture devices to cover potential blind spots of the other image capture device(s).

[0085] Image capture devices 122 may have any suitable resolution capability (e.g., number of pixels associated with the image sensor), and the resolution of the image sensor(s) associated with the image capture device 122 may be higher, lower, or the same as the resolution of the image sensor(s) associated with image capture devices 124 and 126. In some embodiments, the image sensor(s) associated with image capture device 122 and/or image capture devices 124 and 126 may have a resolution of 640 x 480, 1024 x 768, 1280 x 960, or any other suitable resolution.

[0086] The frame rate (e.g., the rate at which an image capture device acquires a set of pixel data of one image frame before moving on to capture pixel data associated with the next image frame) may be controllable. The frame rate associated with image capture device 122 may be higher, lower, or the same as the frame rate associated with image capture devices 124 and 126. The frame rate associated with image capture devices 122, 124, and 126 may depend on a variety of factors that may affect the timing of the frame rate. For example, one or more of image capture devices 122, 124, and 126 may include a selectable pixel delay period imposed before or after acquisition of image data associated with one or more pixels of an image sensor in image capture device 122, 124, and/or 126. Generally, image data corresponding to each pixel may be acquired according to a clock rate for the device (e.g., one pixel per clock cycle). Additionally, in embodiments including a rolling shutter, one or more of image capture devices 122, 124, and 126 may include a selectable horizontal blanking period imposed before or after acquisition of image data associated with a row of pixels of an image sensor in image capture device 122, 124, and/or 126. Further, one or more of image capture devices 122, 124, and/or 126 may include a selectable vertical blanking period imposed before or after acquisition of image data associated with an image frame of image capture device 122, 124, and 126.

[0087] These timing controls may enable synchronization of frame rates associated with image capture devices 122,124, and 126, even where the line scan rates of each are different. Additionally, as will be discussed in greater detail below, these selectable timing controls, among other factors (e.g., image sensor resolution, maximum line scan rates, etc.) may enable synchronization of image capture from an area where the FOV of image capture device 122 overlaps with one or more FOVs of image capture devices 124 and 126, even where the field of view of image capture device 122 is different from the FOVs of image capture devices 124 and 126.

[0088] Frame rate timing in image capture device 122,124, and 126 may depend on the resolution of the associated image sensors. For example, assuming similar line scan rates for both devices, if one device includes an image sensor having a resolution of 640 x 480 and another device includes an image sensor with a resolution of 1280 x 960, then more time will be required to acquire a frame of image data from the sensor having the higher resolution.

[0089] Another factor that may affect the timing of image data acquisition in image capture devices 122, 124, and 126 is the maximum line scan rate. For example, acquisition of a row of image data from an image sensor included in image capture device 122, 124, and 126 will require some minimum amount of time. Assuming no pixel delay periods are added, this minimum amount of time for acquisition of a row of image data will be related to the maximum line scan rate for a particular device. Devices that offer higher maximum line scan rates have the potential to provide higher frame rates than devices with lower maximum line scan rates. In some embodiments, one or more of image capture devices 124 and 126 may have a maximum line scan rate that is higher than a maximum line scan rate associated with image capture device 122. In some embodiments, the maximum line scan rate of image capture device 124 and/or 126 may be 1.25, 1.5, 1.75, or 2 times or more than a maximum line scan rate of image capture device 122.

[0090] In another embodiment, image capture devices 122, 124, and 126 may have the same maximum line scan rate, but image capture device 122 may be operated at a scan rate less than or equal to its maximum scan rate. The system may be configured such that one or more of image capture devices 124 and 126 operate at a line scan rate that is equal to the line scan rate of image capture device 122. In other instances, the system may be configured such that the line scan rate of image capture device 124 and/or image capture device 126 may be 1.25, 1.5, 1.75, or 2 times or more than the line scan rate of image capture device 122.

[0091] In some embodiments, image capture devices 122, 124, and 126 may be asymmetric. That is, they may include cameras having different fields of view (FOV) and focal lengths. The fields of view of image capture devices 122, 124, and 126 may include any desired area relative to an environment of vehicle 200, for example. In some embodiments, one or more of image capture devices 122, 124, and 126 may be configured to acquire image data from an environment in front of vehicle 200, behind vehicle 200, to the sides of vehicle 200, or combinations thereof.

[0092] Further, the focal length associated with each image capture device 122,124, and/or 126 may be selectable (e.g., by inclusion of appropriate lenses etc.) such that each device acquires images of objects at a

IPR2025-00943
Tesla EX1002 Page 1687

desired distance range relative to vehicle 200. For example, in some embodiments image capture devices 122, 124, and 126 may acquire images of close-up objects within a few meters from the vehicle. Image capture devices 122, 124, and 126 may also be configured to acquire images of objects at ranges more distant from the vehicle (e.g., 25 m, 50 m, 100 m, 150 m, or more). Further, the focal lengths of image capture devices 122, 124, and 126 may be selected such that one image capture device (e.g., image capture device 122) can acquire images of objects relatively close to the vehicle (e.g., within 10 m or within 20 m) while the other image capture devices (e.g., image capture devices 124 and 126) can acquire images of more distant objects (e.g., greater than 20 m, 50 m, 100 m, 150 m, etc.) from vehicle 200.

[0093] According to some embodiments, the FOV of one or more image capture devices 122, 124, and 126 may have a wide angle. For example, it may be advantageous to have a FOV of 140 degrees, especially for image capture devices 122, 124, and 126 that may be used to capture images of the area in the vicinity of vehicle 200. For example, image capture device 122 may be used to capture images of the area to the right or left of vehicle 200 and, in such embodiments, it may be desirable for image capture device 122 to have a wide FOV (e.g., at least 140 degrees).

[0094] The field of view associated with each of image capture devices 122, 124, and 126 may depend on the respective focal lengths. For example, as the focal length increases, the corresponding field of view decreases.

[0095] Image capture devices 122, 124, and 126 may be configured to have any suitable fields of view. In one particular example, image capture device 122 may have a horizontal FOV of 46 degrees, image capture device 124 may have a horizontal FOV of 23 degrees, and image capture device 126 may have a horizontal FOV in between 23 and 46 degrees. In another instance, image capture device 122 may have a horizontal FOV of 52 degrees, image capture device 124 may have a horizontal FOV of 26 degrees, and image capture device 126 may have a horizontal FOV in between 26 and 52 degrees. In some embodiments, a ratio of the FOV of image capture device 122 to the FOVs of image capture device 124 and/or image capture device 126 may vary from 1.5 to 2.0. In other embodiments, this ratio may vary between 1.25 and 2.25.

[0096] System 100 may be configured so that a field of view of image capture device 122 overlaps, at least partially or fully, with a field of view of image capture device 124 and/or image capture device 126. In some embodiments, system 100 may be configured such that the fields of view of image capture devices 124 and 126, for example, fall within (e.g., are narrower than) and share a common center with the field of view of image capture device 122. In other embodiments, the image capture devices 122, 124, and 126 may capture adjacent FOVs or may have partial overlap in their FOVs. In some embodiments, the fields of view of image capture devices 122, 124, and 126 may be aligned such that a center of the narrower FOV image capture devices 124 and/or 126 may be located in a lower half of the field of view of the wider FOV device 122.

[0097] FIG. 2F is a diagrammatic representation of exemplary vehicle control systems, consistent with the disclosed embodiments. As indicated in FIG. 2F, vehicle 200 may include throttling system 220, braking system 230, and steering system 240. System 100 may provide inputs (e.g., control signals) to one or more of throttling system 220, braking system 230, and steering system 240 over one or more data links (e.g., any wired and/or wireless link or links for transmitting data). For example, based on analysis of images acquired by image capture devices 122, 124, and/or 126, system 100 may provide control signals to one or more of throttling system 220, braking system 230, and steering system 240 to navigate vehicle 200 (e.g., by causing an acceleration, a turn, a lane shift, etc.). Further, system 100 may receive inputs from one or more of throttling system 220, braking system 230, and steering system 24 indicating operating conditions of vehicle 200 (e.g., speed, whether vehicle 200 is braking and/or turning, etc.). Further details are provided in connection with FIGS. 4-7, below.

[0098] As shown in FIG. 3A, vehicle 200 may also include a user interface 170 for interacting with a driver or a passenger of vehicle 200. For example, user interface 170 in a vehicle application may include a touch screen 320, knobs 330, buttons 340, and a microphone 350. A driver or passenger of vehicle 200 may also use handles (e.g., located on or near the steering column of vehicle 200 including, for example, turn signal handles), buttons (e.g., located on the steering wheel of vehicle 200), and the like, to interact with system 100. In some embodiments, microphone 350 may be positioned adjacent to a rearview mirror 310. Similarly, in some embodiments, image capture device 122 may be located near rearview mirror 310. In some embodiments, user interface 170 may also include one or more speakers 360 (e.g., speakers of a vehicle audio system). For example, system 100 may provide various notifications (e.g., alerts) via speakers 360.

[0099] FIGS. 3B-3D are illustrations of an exemplary camera mount 370 configured to be positioned behind a rearview mirror (e.g., rearview mirror 310) and against a vehicle windshield, consistent with disclosed embodiments. As shown in FIG. 3B, camera mount 370 may include image capture devices 122, 124, and 126. Image capture devices 124 and 126 may be positioned behind a glare shield 380, which may be flush against the vehicle windshield and include a composition of film and/or anti-reflective materials. For example, glare shield 380 may be positioned such that it aligns against a vehicle windshield having a matching slope. In some embodiments, each of image capture devices 122, 124, and 126 may be positioned behind glare shield 380, as depicted, for example, in FIG. 3D. The disclosed embodiments are not limited to any particular configuration of image cap-

IPR2025-00943
Tesla EX1002 Page 1688

ture devices 122, 124, and 126, camera mount 370, and glare shield 380. FIG. 3C is an illustration of camera mount 370 shown in FIG. 3B from a front perspective.

[0100] As will be appreciated by a person skilled in the art having the benefit of this disclosure, numerous variations and/or modifications may be made to the foregoing disclosed embodiments. For example, not all components are essential for the operation of system 100. Further, any component may be located in any appropriate part of system 100 and the components may be rearranged into a variety of configurations while providing the functionality of the disclosed embodiments. Therefore, the foregoing configurations are examples and, regardless of the configurations discussed above, system 100 can provide a wide range of functionality to analyze the surroundings of vehicle 200 and navigate vehicle 200 in response to the analysis.

[0101] As discussed below in further detail and consistent with various disclosed embodiments, system 100 may provide a variety of features related to autonomous driving and/or driver assist technology. For example, system 100 may analyze image data, position data (e.g., GPS location information), map data, speed data, and/or data from sensors included in vehicle 200. System 100 may collect the data for analysis from, for example, image acquisition unit 120, position sensor 130, and other sensors. Further, system 100 may analyze the collected data to determine whether or not vehicle 200 should take a certain action, and then automatically take the determined action without human intervention. For example, when vehicle 200 navigates without human intervention, system 100 may automatically control the braking, acceleration, and/or steering of vehicle 200 (e.g., by sending control signals to one or more of throttling system 220, braking system 230, and steering system 240). Further, system 100 may analyze the collected data and issue warnings and/or alerts to vehicle occupants based on the analysis of the collected data.

[0102] Further, consistent with disclosed embodiments, the functionality provided by system 100 may cause vehicle 200 to take different actions to navigate vehicle 200 within a lane and/or relative to other vehicles and/or objects. For example, system 100 may adjust the positioning of vehicle 200 relative to a lane within which vehicle 200 is traveling and/or relative to objects positioned near vehicle 200, select a particular lane for vehicle 200 to use while traveling, and take action in response to an encroaching vehicle, such as a vehicle attempting to move into the lane within which vehicle 200 is traveling. Additionally, system 100 may control the speed of vehicle 200 in different scenarios, such as when vehicle 200 is making a turn. System 100 may further cause vehicle 200 to mimic the actions of a leading vehicle or monitor a target vehicle and navigate vehicle 200 so that it passes the target vehicle. Additional details regarding the various embodiments that are provided by system 100 are provided below.

Forward-Facing Multi-Imaging System

[0103] As discussed above, system 100 may provide drive assist functionality that uses a multi-camera system. The multi-camera system may use one or more cameras facing in the forward direction of a vehicle. In other embodiments, the multi-camera system may include one or more cameras facing to the side of a vehicle or to the rear of the vehicle. In one embodiment, for example, system 100 may use a two-camera imaging system, where a first camera and a second camera (e.g., image capture devices 122 and 124) may be positioned at the front and/or the sides of a vehicle (e.g., vehicle 200). The first camera may have a field of view that is greater than, less than, or partially overlapping with, the field of view of the second camera. In addition, the first camera may be connected to a first image processor to perform monocular image analysis of images provided by the first camera, and the second camera may be connected to a second image processor to perform monocular image analysis of images provided by the second camera. The outputs (e.g., processed information) of the first and second image processors may be combined. In some embodiments, the second image processor may receive images from both the first camera and second camera to perform stereo analysis. In another embodiment, system 100 may use a three-camera imaging system where each of the cameras has a different field of view. Such a system may, therefore, make decisions based on information derived from objects located at varying distances both forward and to the sides of the vehicle. References to monocular image analysis may refer to instances where image analysis is performed based on images captured from a single point of view (e.g., from a single camera). Stereo image analysis may refer to instances where image analysis is performed based on two or more images captured with one or more variations of an image capture parameter. For example, captured images suitable for performing stereo image analysis may include images captured: from two or more different positions, from different fields of view, using different focal lengths, along with parallax information, etc.

[0104] For example, in one embodiment, system 100 may implement a three camera configuration using image capture devices 122-126. In such a configuration, image capture device 122 may provide a narrow field of view (e.g., 34 degrees, or other values selected from a range of about 20 to 45 degrees, etc.), image capture device 124 may provide a wide field of view (e.g., 150 degrees or other values selected from a range of about 100 to about 180 degrees), and image capture device 126 may provide an intermediate field of view (e.g., 46 degrees or other values selected from a range of about 35 to about 60 degrees). In some embodiments, image capture device 126 may act as a main or primary camera. Image capture devices 122-126 may be positioned behind rearview mirror 310 and positioned substantially side-by-side (e.g., 6 cm apart). Further, in some embod-

IPR2025-00943
Tesla EX1002 Page 1689

iments, as discussed above, one or more of image capture devices 122-126 may be mounted behind glare shield 380 that is flush with the windshield of vehicle 200. Such shielding may act to minimize the impact of any reflections from inside the car on image capture devices 122-126.

[0105] In another embodiment, as discussed above in connection with FIGS. 3B and 3C, the wide field of view camera (e.g., image capture device 124 in the above example) may be mounted lower than the narrow and main field of view cameras (e.g., image devices 122 and 126 in the above example). This configuration may provide a free line of sight from the wide field of view camera. To reduce reflections, the cameras may be mounted close to the windshield of vehicle 200, and may include polarizers on the cameras to damp reflected light.

[0106] A three camera system may provide certain performance characteristics. For example, some embodiments may include an ability to validate the detection of objects by one camera based on detection results from another camera. In the three camera configuration discussed above, processing unit 110 may include, for example, three processing devices (e.g., three EyeQ series of processor chips, as discussed above), with each processing device dedicated to processing images captured by one or more of image capture devices 122-126.

[0107] In a three camera system, a first processing device may receive images from both the main camera and the narrow field of view camera, and perform vision processing of the narrow FOV camera to, for example, detect other vehicles, pedestrians, lane marks, traffic signs, traffic lights, and other road objects. Further, the first processing device may calculate a disparity of pixels between the images from the main camera and the narrow camera and create a 3D reconstruction of the environment of vehicle 200. The first processing device may then combine the 3D reconstruction with 3D map data or with 3D information calculated based on information from another camera.

[0108] The second processing device may receive images from main camera and perform vision processing to detect other vehicles, pedestrians, lane marks, traffic signs, traffic lights, and other road objects. Additionally, the second processing device may calculate a camera displacement and, based on the displacement, calculate a disparity of pixels between successive images and create a 3D reconstruction of the scene (e.g., a structure from motion). The second processing device may send the structure from motion based 3D reconstruction to the first processing device to be combined with the stereo 3D images.

[0109] The third processing device may receive images from the wide FOV camera and process the images to detect vehicles, pedestrians, lane marks, traffic signs, traffic lights, and other road objects. The third processing device may further execute additional processing instructions to analyze images to identify objects moving in the image, such as vehicles changing lanes, pedestrians, etc.

[0110] In some embodiments, having streams of image-based information captured and processed independently may provide an opportunity for providing redundancy in the system. Such redundancy may include, for example, using a first image capture device and the images processed from that device to validate and/or supplement information obtained by capturing and processing image information from at least a second image capture device.

[0111] In some embodiments, system 100 may use two image capture devices (e.g., image capture devices 122 and 124) in providing navigation assistance for vehicle 200 and use a third image capture device (e.g., image capture device 126) to provide redundancy and validate the analysis of data received from the other two image capture devices. For example, in such a configuration, image capture devices 122 and 124 may provide images for stereo analysis by system 100 for navigating vehicle 200, while image capture device 126 may provide images for monocular analysis by system 100 to provide redundancy and validation of information obtained based on images captured from image capture device 122 and/or image capture device 124. That is, image capture device 126 (and a corresponding processing device) may be considered to provide a redundant sub-system for providing a check on the analysis derived from image capture devices 122 and 124 (e.g., to provide an automatic emergency braking (AEB) system).

[0112] One of skill in the art will recognize that the above camera configurations, camera placements, number of cameras, camera locations, etc., are examples only. These components and others described relative to the overall system may be assembled and used in a variety of different configurations without departing from the scope of the disclosed embodiments. Further details regarding usage of a multi-camera system to provide driver assist and/or autonomous vehicle functionality follow below.

[0113] FIG. 4 is an exemplary functional block diagram of memory 140 and/or 150, which may be stored/programmed with instructions for performing one or more operations consistent with the disclosed embodiments. Although the following refers to memory 140, one of skill in the art will recognize that instructions may be stored in memory 140 and/or 150.

[0114] As shown in FIG. 4, memory 140 may store a monocular image analysis module 402, a stereo image analysis module 404, a velocity and acceleration module 406, and a navigational response module 408. The disclosed embodiments are not limited to any particular configuration of memory 140. Further, application processor 180 and/or image processor 190 may execute the instructions stored in any of modules 402-408 included in memory 140. One of skill in the art will understand that references in the following discussions to processing unit 110 may refer to application processor 180 and image processor 190 individually or collectively. Accordingly,

IPR2025-00943
Tesla EX1002 Page 1690

steps of any of the following processes may be performed by one or more processing devices.

[0115] In one embodiment, monocular image analysis module 402 may store instructions (such as computer vision software) which, when executed by processing unit 110, performs monocular image analysis of a set of images acquired by one of image capture devices 122, 124, and 126. In some embodiments, processing unit 110 may combine information from a set of images with additional sensory information (e.g., information from radar) to perform the monocular image analysis. As described in connection with FIGS. 5A-5D below, monocular image analysis module 402 may include instructions for detecting a set of features within the set of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, hazardous objects, and any other feature associated with an environment of a vehicle. Based on the analysis, system 100 (e.g., via processing unit 110) may cause one or more navigational responses in vehicle 200, such as a turn, a lane shift, a change in acceleration, and the like, as discussed below in connection with navigational response module 408.

[0116] In one embodiment, stereo image analysis module 404 may store instructions (such as computer vision software) which, when executed by processing unit 110, performs stereo image analysis of first and second sets of images acquired by a combination of image capture devices selected from any of image capture devices 122, 124, and 126. In some embodiments, processing unit 110 may combine information from the first and second sets of images with additional sensory information (e.g., information from radar) to perform the stereo image analysis. For example, stereo image analysis module 404 may include instructions for performing stereo image analysis based on a first set of images acquired by image capture device 124 and a second set of images acquired by image capture device 126. As described in connection with FIG. 6 below, stereo image analysis module 404 may include instructions for detecting a set of features within the first and second sets of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, hazardous objects, and the like. Based on the analysis, processing unit 110 may cause one or more navigational responses in vehicle 200, such as a turn, a lane shift, a change in acceleration, and the like, as discussed below in connection with navigational response module 408.

[0117] In one embodiment, velocity and acceleration module 406 may store software configured to analyze data received from one or more computing and electromechanical devices in vehicle 200 that are configured to cause a change in velocity and/or acceleration of vehicle 200. For example, processing unit 110 may execute instructions associated with velocity and acceleration module 406 to calculate a target speed for vehicle 200 based on data derived from execution of monocular image analysis module 402 and/or stereo image analysis module 404. Such data may include, for example, a target position, velocity, and/or acceleration, the position and/or speed of vehicle 200 relative to a nearby vehicle, pedestrian, or road object, position information for vehicle 200 relative to lane markings of the road, and the like. In addition, processing unit 110 may calculate a target speed for vehicle 200 based on sensory input (e.g., information from radar) and input from other systems of vehicle 200, such as throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. Based on the calculated target speed, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to trigger a change in velocity and/or acceleration by, for example, physically depressing the brake or easing up off the accelerator of vehicle 200.

[0118] In one embodiment, navigational response module 408 may store software executable by processing unit 110 to determine a desired navigational response based on data derived from execution of monocular image analysis module 402 and/or stereo image analysis module 404. Such data may include position and speed information associated with nearby vehicles, pedestrians, and road objects, target position information for vehicle 200, and the like. Additionally, in some embodiments, the navigational response may be based (partially or fully) on map data, a predetermined position of vehicle 200, and/or a relative velocity or a relative acceleration between vehicle 200 and one or more objects detected from execution of monocular image analysis module 402 and/or stereo image analysis module 404. Navigational response module 408 may also determine a desired navigational response based on sensory input (e.g., information from radar) and inputs from other systems of vehicle 200, such as throttling system 220, braking system 230, and steering system 240 of vehicle 200. Based on the desired navigational response, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and steering system 240 of vehicle 200 to trigger a desired navigational response by, for example, turning the steering wheel of vehicle 200 to achieve a rotation of a predetermined angle. In some embodiments, processing unit 110 may use the output of navigational response module 408 (e.g., the desired navigational response) as an input to execution of velocity and acceleration module 406 for calculating a change in speed of vehicle 200.

[0119] FIG. 5A is a flowchart showing an exemplary process 500A for causing one or more navigational responses based on monocular image analysis, consistent with disclosed embodiments. At step 510, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120. For instance, a camera included in image acquisition unit 120 (such as image capture device 122 having field of view 202) may capture a plurality of images of an area forward of vehicle 200 (or to the sides or rear of a vehicle, for example) and transmit them over a data connection (e.g., digital, wired, USB, wireless, Bluetooth,

IPR2025-00943
Tesla EX1002 Page 1691

etc.) to processing unit 110. Processing unit 110 may execute monocular image analysis module 402 to analyze the plurality of images at step 520, as described in further detail in connection with FIGS. 5B-5D below. By performing the analysis, processing unit 110 may detect a set of features within the set of images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, and the like.

[0120] Processing unit 110 may also execute monocular image analysis module 402 to detect various road hazards at step 520, such as, for example, parts of a truck tire, fallen road signs, loose cargo, small animals, and the like. Road hazards may vary in structure, shape, size, and color, which may make detection of such hazards more challenging. In some embodiments, processing unit 110 may execute monocular image analysis module 402 to perform multi-frame analysis on the plurality of images to detect road hazards. For example, processing unit 110 may estimate camera motion between consecutive image frames and calculate the disparities in pixels between the frames to construct a 3D-map of the road. Processing unit 110 may then use the 3D-map to detect the road surface, as well as hazards existing above the road surface.

[0121] At step 530, processing unit 110 may execute navigational response module 408 to cause one or more navigational responses in vehicle 200 based on the analysis performed at step 520 and the techniques as described above in connection with FIG. 4. Navigational responses may include, for example, a turn, a lane shift, a change in acceleration, and the like. In some embodiments, processing unit 110 may use data derived from execution of velocity and acceleration module 406 to cause the one or more navigational responses. Additionally, multiple navigational responses may occur simultaneously, in sequence, or any combination thereof. For instance, processing unit 110 may cause vehicle 200 to shift one lane over and then accelerate by, for example, sequentially transmitting control signals to steering system 240 and throttling system 220 of vehicle 200. Alternatively, processing unit 110 may cause vehicle 200 to brake while at the same time shifting lanes by, for example, simultaneously transmitting control signals to braking system 230 and steering system 240 of vehicle 200.

[0122] FIG. 5B is a flowchart showing an exemplary process 500B for detecting one or more vehicles and/or pedestrians in a set of images, consistent with disclosed embodiments. Processing unit 110 may execute monocular image analysis module 402 to implement process 500B. At step 540, processing unit 110 may determine a set of candidate objects representing possible vehicles and/or pedestrians. For example, processing unit 110 may scan one or more images, compare the images to one or more predetermined patterns, and identify within each image possible locations that may contain objects of interest (e.g., vehicles, pedestrians, or portions thereof). The predetermined patterns may be designed in such a way to achieve a high rate of "false hits" and a low rate of "misses." For example, processing unit 110 may use a low threshold of similarity to predetermined patterns for identifying candidate objects as possible vehicles or pedestrians. Doing so may allow processing unit 110 to reduce the probability of missing (e.g., not identifying) a candidate object representing a vehicle or pedestrian.

[0123] At step 542, processing unit 110 may filter the set of candidate objects to exclude certain candidates (e.g., irrelevant or less relevant objects) based on classification criteria. Such criteria may be derived from various properties associated with object types stored in a database (e.g., a database stored in memory 140). Properties may include object shape, dimensions, texture, position (e.g., relative to vehicle 200), and the like. Thus, processing unit 110 may use one or more sets of criteria to reject false candidates from the set of candidate objects.

[0124] At step 544, processing unit 110 may analyze multiple frames of images to determine whether objects in the set of candidate objects represent vehicles and/or pedestrians. For example, processing unit 110 may track a detected candidate object across consecutive frames and accumulate frame-by-frame data associated with the detected object (e.g., size, position relative to vehicle 200, etc.). Additionally, processing unit 110 may estimate parameters for the detected object and compare the object's frame-by-frame position data to a predicted position.

[0125] At step 546, processing unit 110 may construct a set of measurements for the detected objects. Such measurements may include, for example, position, velocity, and acceleration values (relative to vehicle 200) associated with the detected objects. In some embodiments, processing unit 110 may construct the measurements based on estimation techniques using a series of time-based observations such as Kalman filters or linear quadratic estimation (LQE), and/or based on available modeling data for different object types (e.g., cars, trucks, pedestrians, bicycles, road signs, etc.). The Kalman filters may be based on a measurement of an object's scale, where the scale measurement is proportional to a time to collision (e.g., the amount of time for vehicle 200 to reach the object). Thus, by performing steps 540-546, processing unit 110 may identify vehicles and pedestrians appearing within the set of captured images and derive information (e.g., position, speed, size) associated with the vehicles and pedestrians. Based on the identification and the derived information, processing unit 110 may cause one or more navigational responses in vehicle 200, as described in connection with FIG. 5A, above.

[0126] At step 548, processing unit 110 may perform an optical flow analysis of one or more images to reduce the probabilities of detecting a "false hit" and missing a candidate object that represents a vehicle or pedestrian. The optical flow analysis may refer to, for example, analyzing motion patterns relative to vehicle 200 in the one or more images associated with other vehicles and pedestrians, and that are distinct from road surface motion.

IPR2025-00943
Tesla EX1002 Page 1692

Processing unit 110 may calculate the motion of candidate objects by observing the different positions of the objects across multiple image frames, which are captured at different times. Processing unit 110 may use the position and time values as inputs into mathematical models for calculating the motion of the candidate objects. Thus, optical flow analysis may provide another method of detecting vehicles and pedestrians that are nearby vehicle 200. Processing unit 110 may perform optical flow analysis in combination with steps 540-546 to provide redundancy for detecting vehicles and pedestrians and increase the reliability of system 100.

[0127] FIG. 5C is a flowchart showing an exemplary process 500C for detecting road marks and/or lane geometry information in a set of images, consistent with disclosed embodiments. Processing unit 110 may execute monocular image analysis module 402 to implement process 500C. At step 550, processing unit 110 may detect a set of objects by scanning one or more images. To detect segments of lane markings, lane geometry information, and other pertinent road marks, processing unit 110 may filter the set of objects to exclude those determined to be irrelevant (e.g., minor potholes, small rocks, etc.). At step 552, processing unit 110 may group together the segments detected in step 550 belonging to the same road mark or lane mark. Based on the grouping, processing unit 110 may develop a model to represent the detected segments, such as a mathematical model.

[0128] At step 554, processing unit 110 may construct a set of measurements associated with the detected segments. In some embodiments, processing unit 110 may create a projection of the detected segments from the image plane onto the real-world plane. The projection may be characterized using a 3rd-degree polynomial having coefficients corresponding to physical properties such as the position, slope, curvature, and curvature derivative of the detected road. In generating the projection, processing unit 110 may take into account changes in the road surface, as well as pitch and roll rates associated with vehicle 200. In addition, processing unit 110 may model the road elevation by analyzing position and motion cues present on the road surface. Further, processing unit 110 may estimate the pitch and roll rates associated with vehicle 200 by tracking a set of feature points in the one or more images.

[0129] At step 556, processing unit 110 may perform multi-frame analysis by, for example, tracking the detected segments across consecutive image frames and accumulating frame-by-frame data associated with detected segments. As processing unit 110 performs multi-frame analysis, the set of measurements constructed at step 554 may become more reliable and associated with an increasingly higher confidence level. Thus, by performing steps 550-556, processing unit 110 may identify road marks appearing within the set of captured images and derive lane geometry information. Based on the identification and the derived information, processing unit 110 may cause one or more navigational responses in vehicle 200, as described in connection with FIG. 5A, above.

[0130] At step 558, processing unit 110 may consider additional sources of information to further develop a safety model for vehicle 200 in the context of its surroundings. Processing unit 110 may use the safety model to define a context in which system 100 may execute autonomous control of vehicle 200 in a safe manner. To develop the safety model, in some embodiments, processing unit 110 may consider the position and motion of other vehicles, the detected road edges and barriers, and/or general road shape descriptions extracted from map data (such as data from map database 160). By considering additional sources of information, processing unit 110 may provide redundancy for detecting road marks and lane geometry and increase the reliability of system 100.

[0131] FIG. 5D is a flowchart showing an exemplary process 500D for detecting traffic lights in a set of images, consistent with disclosed embodiments. Processing unit 110 may execute monocular image analysis module 402 to implement process 500D. At step 560, processing unit 110 may scan the set of images and identify objects appearing at locations in the images likely to contain traffic lights. For example, processing unit 110 may filter the identified objects to construct a set of candidate objects, excluding those objects unlikely to correspond to traffic lights. The filtering may be done based on various properties associated with traffic lights, such as shape, dimensions, texture, position (e.g., relative to vehicle 200), and the like. Such properties may be based on multiple examples of traffic lights and traffic control signals and stored in a database. In some embodiments, processing unit 110 may perform multi-frame analysis on the set of candidate objects reflecting possible traffic lights. For example, processing unit 110 may track the candidate objects across consecutive image frames, estimate the real-world position of the candidate objects, and filter out those objects that are moving (which are unlikely to be traffic lights). In some embodiments, processing unit 110 may perform color analysis on the candidate objects and identify the relative position of the detected colors appearing inside possible traffic lights.

[0132] At step 562, processing unit 110 may analyze the geometry of a junction. The analysis may be based on any combination of: (i) the number of lanes detected on either side of vehicle 200, (ii) markings (such as arrow marks) detected on the road, and (iii) descriptions of the junction extracted from map data (such as data from map database 160). Processing unit 110 may conduct the analysis using information derived from execution of monocular analysis module 402. In addition, Processing unit 110 may determine a correspondence between the traffic lights detected at step 560 and the lanes appearing near vehicle 200.

[0133] As vehicle 200 approaches the junction, at step 564, processing unit 110 may update the confidence level associated with the analyzed junction geometry and the detected traffic lights. For instance, the number of traffic

IPR2025-00943
Tesla EX1002 Page 1693

lights estimated to appear at the junction as compared with the number actually appearing at the junction may impact the confidence level. Thus, based on the confidence level, processing unit 110 may delegate control to the driver of vehicle 200 in order to improve safety conditions. By performing steps 560-564, processing unit 110 may identify traffic lights appearing within the set of captured images and analyze junction geometry information. Based on the identification and the analysis, processing unit 110 may cause one or more navigational responses in vehicle 200, as described in connection with FIG. 5A, above.

[0134] FIG. 5E is a flowchart showing an exemplary process 500E for causing one or more navigational responses in vehicle 200 based on a vehicle path, consistent with the disclosed embodiments. At step 570, processing unit 110 may construct an initial vehicle path associated with vehicle 200. The vehicle path may be represented using a set of points expressed in coordinates (x, z), and the distance d_i between two points in the set of points may fall in the range of 1 to 5 meters. In one embodiment, processing unit 110 may construct the initial vehicle path using two polynomials, such as left and right road polynomials. Processing unit 110 may calculate the geometric midpoint between the two polynomials and offset each point included in the resultant vehicle path by a predetermined offset (e.g., a smart lane offset), if any (an offset of zero may correspond to travel in the middle of a lane). The offset may be in a direction perpendicular to a segment between any two points in the vehicle path. In another embodiment, processing unit 110 may use one polynomial and an estimated lane width to offset each point of the vehicle path by half the estimated lane width plus a predetermined offset (e.g., a smart lane offset).

[0135] At step 572, processing unit 110 may update the vehicle path constructed at step 570. Processing unit 110 may reconstruct the vehicle path constructed at step 570 using a higher resolution, such that the distance d_k between two points in the set of points representing the vehicle path is less than the distance d_i described above. For example, the distance d_k may fall in the range of 0.1 to 0.3 meters. Processing unit 110 may reconstruct the vehicle path using a parabolic spline algorithm, which may yield a cumulative distance vector S corresponding to the total length of the vehicle path (i.e., based on the set of points representing the vehicle path).

[0136] At step 574, processing unit 110 may determine a look-ahead point (expressed in coordinates as (x_l, z_l)) based on the updated vehicle path constructed at step 572. Processing unit 110 may extract the look-ahead point from the cumulative distance vector S, and the look-ahead point may be associated with a look-ahead distance and look-ahead time. The look-ahead distance, which may have a lower bound ranging from 10 to 20 meters, may be calculated as the product of the speed of vehicle 200 and the look-ahead time. For example, as the speed of vehicle 200 decreases, the look-ahead distance may also decrease (e.g., until it reaches the lower bound). The look-ahead time, which may range from 0.5 to 1.5 seconds, may be inversely proportional to the gain of one or more control loops associated with causing a navigational response in vehicle 200, such as the heading error tracking control loop. For example, the gain of the heading error tracking control loop may depend on the bandwidth of a yaw rate loop, a steering actuator loop, car lateral dynamics, and the like. Thus, the higher the gain of the heading error tracking control loop, the lower the look-ahead time.

[0137] At step 576, processing unit 110 may determine a heading error and yaw rate command based on the look-ahead point determined at step 574. Processing unit 110 may determine the heading error by calculating the arctangent of the look-ahead point, e.g., arctan (x_l/z_l). Processing unit 110 may determine the yaw rate command as the product of the heading error and a high-level control gain. The high-level control gain may be equal to: (2 / look-ahead time), if the look-ahead distance is not at the lower bound. Otherwise, the high-level control gain may be equal to: (2 * speed of vehicle 200 / look-ahead distance).

[0138] FIG. 5F is a flowchart showing an exemplary process 500F for determining whether a leading vehicle is changing lanes, consistent with the disclosed embodiments. At step 580, processing unit 110 may determine navigation information associated with a leading vehicle (e.g., a vehicle traveling ahead of vehicle 200). For example, processing unit 110 may determine the position, velocity (e.g., direction and speed), and/or acceleration of the leading vehicle, using the techniques described in connection with FIGS. 5A and 5B, above. Processing unit 110 may also determine one or more road polynomials, a look-ahead point (associated with vehicle 200), and/or a snail trail (e.g., a set of points describing a path taken by the leading vehicle), using the techniques described in connection with FIG. 5E, above.

[0139] At step 582, processing unit 110 may analyze the navigation information determined at step 580. In one embodiment, processing unit 110 may calculate the distance between a snail trail and a road polynomial (e.g., along the trail). If the variance of this distance along the trail exceeds a predetermined threshold (for example, 0.1 to 0.2 meters on a straight road, 0.3 to 0.4 meters on a moderately curvy road, and 0.5 to 0.6 meters on a road with sharp curves), processing unit 110 may determine that the leading vehicle is likely changing lanes. In the case where multiple vehicles are detected traveling ahead of vehicle 200, processing unit 110 may compare the snail trails associated with each vehicle. Based on the comparison, processing unit 110 may determine that a vehicle whose snail trail does not match with the snail trails of the other vehicles is likely changing lanes. Processing unit 110 may additionally compare the curvature of the snail trail (associated with the leading vehicle) with the expected curvature of the road segment in which the leading vehicle is traveling. The expected

IPR2025-00943
Tesla EX1002 Page 1694

curvature may be extracted from map data (e.g., data from map database 160), from road polynomials, from other vehicles' snail trails, from prior knowledge about the road, and the like. If the difference in curvature of the snail trail and the expected curvature of the road segment exceeds a predetermined threshold, processing unit 110 may determine that the leading vehicle is likely changing lanes.

[0140] In another embodiment, processing unit 110 may compare the leading vehicle's instantaneous position with the look-ahead point (associated with vehicle 200) over a specific period of time (e.g., 0.5 to 1.5 seconds). If the distance between the leading vehicle's instantaneous position and the look-ahead point varies during the specific period of time, and the cumulative sum of variation exceeds a predetermined threshold (for example, 0.3 to 0.4 meters on a straight road, 0.7 to 0.8 meters on a moderately curvy road, and 1.3 to 1.7 meters on a road with sharp curves), processing unit 110 may determine that the leading vehicle is likely changing lanes. In another embodiment, processing unit 110 may analyze the geometry of the snail trail by comparing the lateral distance traveled along the trail with the expected curvature of the snail trail. The expected radius of curvature may be determined according to the calculation: $(\delta_x^2 + \delta_x^2) / 2 / (\delta_x)$, where δ_x represents the lateral distance traveled and δ_z represents the longitudinal distance traveled. If the difference between the lateral distance traveled and the expected curvature exceeds a predetermined threshold (e.g., 500 to 700 meters), processing unit 110 may determine that the leading vehicle is likely changing lanes. In another embodiment, processing unit 110 may analyze the position of the leading vehicle. If the position of the leading vehicle obscures a road polynomial (e.g., the leading vehicle is overlaid on top of the road polynomial), then processing unit 110 may determine that the leading vehicle is likely changing lanes. In the case where the position of the leading vehicle is such that, another vehicle is detected ahead of the leading vehicle and the snail trails of the two vehicles are not parallel, processing unit 110 may determine that the (closer) leading vehicle is likely changing lanes.

[0141] At step 584, processing unit 110 may determine whether or not leading vehicle 200 is changing lanes based on the analysis performed at step 582. For example, processing unit 110 may make the determination based on a weighted average of the individual analyses performed at step 582. Under such a scheme, for example, a decision by processing unit 110 that the leading vehicle is likely changing lanes based on a particular type of analysis may be assigned a value of "1" (and "0" to represent a determination that the leading vehicle is not likely changing lanes). Different analyses performed at step 582 may be assigned different weights, and the disclosed embodiments are not limited to any particular combination of analyses and weights.

[0142] FIG. 6 is a flowchart showing an exemplary process 600 for causing one or more navigational responses based on stereo image analysis, consistent with disclosed embodiments. At step 610, processing unit 110 may receive a first and second plurality of images via data interface 128. For example, cameras included in image acquisition unit 120 (such as image capture devices 122 and 124 having fields of view 202 and 204) may capture a first and second plurality of images of an area forward of vehicle 200 and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive the first and second plurality of images via two or more data interfaces. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0143] At step 620, processing unit 110 may execute stereo image analysis module 404 to perform stereo image analysis of the first and second plurality of images to create a 3D map of the road in front of the vehicle and detect features within the images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, road hazards, and the like. Stereo image analysis may be performed in a manner similar to the steps described in connection with FIGS. 5A-5D, above. For example, processing unit 110 may execute stereo image analysis module 404 to detect candidate objects (e.g., vehicles, pedestrians, road marks, traffic lights, road hazards, etc.) within the first and second plurality of images, filter out a subset of the candidate objects based on various criteria, and perform multi-frame analysis, construct measurements, and determine a confidence level for the remaining candidate objects. In performing the steps above, processing unit 110 may consider information from both the first and second plurality of images, rather than information from one set of images alone. For example, processing unit 110 may analyze the differences in pixel-level data (or other data subsets from among the two streams of captured images) for a candidate object appearing in both the first and second plurality of images. As another example, processing unit 110 may estimate a position and/or velocity of a candidate object (e.g., relative to vehicle 200) by observing that the object appears in one of the plurality of images but not the other or relative to other differences that may exist relative to objects appearing if the two image streams. For example, position, velocity, and/or acceleration relative to vehicle 200 may be determined based on trajectories, positions, movement characteristics, etc. of features associated with an object appearing in one or both of the image streams.

[0144] At step 630, processing unit 110 may execute navigational response module 408 to cause one or more navigational responses in vehicle 200 based on the analysis performed at step 620 and the techniques as described above in connection with FIG. 4. Navigational responses may include, for example, a turn, a lane shift, a change in acceleration, a change in velocity, braking, and the like. In some embodiments, processing unit 110 may use data derived from execution of velocity and ac-

IPR2025-00943
Tesla EX1002 Page 1695

celeration module 406 to cause the one or more navigational responses. Additionally, multiple navigational responses may occur simultaneously, in sequence, or any combination thereof.

[0145] FIG. 7 is a flowchart showing an exemplary process 700 for causing one or more navigational responses based on an analysis of three sets of images, consistent with disclosed embodiments. At step 710, processing unit 110 may receive a first, second, and third plurality of images via data interface 128. For instance, cameras included in image acquisition unit 120 (such as image capture devices 122, 124, and 126 having fields of view 202, 204, and 206) may capture a first, second, and third plurality of images of an area forward and/or to the side of vehicle 200 and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive the first, second, and third plurality of images via three or more data interfaces. For example, each of image capture devices 122, 124, 126 may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0146] At step 720, processing unit 110 may analyze the first, second, and third plurality of images to detect features within the images, such as lane markings, vehicles, pedestrians, road signs, highway exit ramps, traffic lights, road hazards, and the like. The analysis may be performed in a manner similar to the steps described in connection with FIGS. 5A-5D and 6, above. For instance, processing unit 110 may perform monocular image analysis (e.g., via execution of monocular image analysis module 402 and based on the steps described in connection with FIGS. 5A-5D, above) on each of the first, second, and third plurality of images. Alternatively, processing unit 110 may perform stereo image analysis (e.g., via execution of stereo image analysis module 404 and based on the steps described in connection with FIG. 6, above) on the first and second plurality of images, the second and third plurality of images, and/or the first and third plurality of images. The processed information corresponding to the analysis of the first, second, and/or third plurality of images may be combined. In some embodiments, processing unit 110 may perform a combination of monocular and stereo image analyses. For example, processing unit 110 may perform monocular image analysis (e.g., via execution of monocular image analysis module 402) on the first plurality of images and stereo image analysis (e.g., via execution of stereo image analysis module 404) on the second and third plurality of images. The configuration of image capture devices 122, 124, and 126-including their respective locations and fields of view 202, 204, and 206-may influence the types of analyses conducted on the first, second, and third plurality of images. The disclosed embodiments are not limited to a particular configuration of image capture devices 122, 124, and 126, or the types of analyses conducted on the first, second, and third plurality of images.

[0147] In some embodiments, processing unit 110 may perform testing on system 100 based on the images acquired and analyzed at steps 710 and 720. Such testing may provide an indicator of the overall performance of system 100 for certain configurations of image capture devices 122, 124, and 126. For example, processing unit 110 may determine the proportion of "false hits" (e.g., cases where system 100 incorrectly determined the presence of a vehicle or pedestrian) and "misses."

[0148] At step 730, processing unit 110 may cause one or more navigational responses in vehicle 200 based on information derived from two of the first, second, and third plurality of images. Selection of two of the first, second, and third plurality of images may depend on various factors, such as, for example, the number, types, and sizes of objects detected in each of the plurality of images. Processing unit 110 may also make the selection based on image quality and resolution, the effective field of view reflected in the images, the number of captured frames, the extent to which one or more objects of interest actually appear in the frames (e.g., the percentage of frames in which an object appears, the proportion of the object that appears in each such frame, etc.), and the like.

[0149] In some embodiments, processing unit 110 may select information derived from two of the first, second, and third plurality of images by determining the extent to which information derived from one image source is consistent with information derived from other image sources. For example, processing unit 110 may combine the processed information derived from each of image capture devices 122, 124, and 126 (whether by monocular analysis, stereo analysis, or any combination of the two) and determine visual indicators (e.g., lane markings, a detected vehicle and its location and/or path, a detected traffic light, etc.) that are consistent across the images captured from each of image capture devices 122, 124, and 126. Processing unit 110 may also exclude information that is inconsistent across the captured images (e.g., a vehicle changing lanes, a lane model indicating a vehicle that is too close to vehicle 200, etc.). Thus, processing unit 110 may select information derived from two of the first, second, and third plurality of images based on the determinations of consistent and inconsistent information.

[0150] Navigational responses may include, for example, a turn, a lane shift, a change in acceleration, and the like. Processing unit 110 may cause the one or more navigational responses based on the analysis performed at step 720 and the techniques as described above in connection with FIG. 4. Processing unit 110 may also use data derived from execution of velocity and acceleration module 406 to cause the one or more navigational responses. In some embodiments, processing unit 110 may cause the one or more navigational responses based on a relative position, relative velocity, and/or relative acceleration between vehicle 200 and an object detected within any of the first, second, and third plurality

IPR2025-00943
Tesla EX1002 Page 1696

of images. Multiple navigational responses may occur simultaneously, in sequence, or any combination thereof.

Smart Lane Offset

[0151] System 100 may provide driver assist functionality that monitors and adjusts the free space between vehicle 200 and boundaries (e.g., lane boundaries) within which vehicle 200 is traveling. A lane may refer to a designated or intended travel path of a vehicle and may have marked (e.g., lines on a road) or unmarked boundaries (e.g., an edge of a road, a road barrier, guard rail, parked vehicles, etc.). For instance, as a default, system 100 may maximize or increase the free space between vehicle 200 and the current lane boundaries. Further, under certain conditions, system 100 may offset vehicle 200 toward one side of the current lane of travel without leaving the lane. For example, where vehicle 200 is traveling around a curve, system 100 may decrease the offset towards the inside of the curve (e.g., adjust the position of vehicle 200 so that it is closer to the inside of the curve). As another example, where obstacles-such as parked cars, pedestrians, or cyclists-are present on one side of the lane, system 100 may increase the offset on the side of vehicle 200 where the obstacles are located (e.g., adjust the position of vehicle 200 so that it is further from the obstacles).

[0152] FIG. 8A illustrates vehicle 200 traveling on a roadway 800 in which the disclosed systems and methods for identifying lane constraints and operating vehicle 200 within the lane constraints may be used. As shown, vehicle 200 may have a first vehicle side 802, which may be a first distance 805 from a first lane constraint 810. Similarly, vehicle 200 may have a second vehicle side 812 opposite from first vehicle side 802, and second vehicle side 812 may be a second distance 815 from second lane constraint 820. In this manner, first lane constraint 810 and second lane constraint 820 may define a lane within which vehicle 200 may travel.

[0153] Processing unit 110 may be configured to determine first lane constraint 810 and second lane constraint 820 based on a plurality of images acquired by image capture device 122-126 that processing unit 110 may receive via data interface 128. According to some embodiments, first lane constraint 810 and/or second lane constraint 820 may be identified by visible lane boundaries, such as dashed or solid lines marked on a road surface. Additionally or alternatively, first lane constraint 810 and/or second lane constraint 820 may include an edge of a road surface or a barrier. Additionally or alternatively, first lane constraint 810 and/or second lane constrain 820 may include markers (e.g., Botts' dots). According to some embodiments, processing unit 110 may determine first lane constraint 810 and/or second lane constraint 820 by identifying a midpoint of a road surface width. According to some embodiments, if processing unit 110 identifies only first lane constraint 810, processing unit 110 may estimate second lane constraint 820, such as based on an estimated lane width or road width. Processing unit 110 may identify lane constraints in this manner when, for example, lines designating road lanes are not painted or otherwise labeled.

[0154] Detection of first lane constraint 810 and/or second lane constraint 820 may include processing unit 110 determining their 3D models in camera coordinate system. For example, the 3D models of first lane constraint 810 and/or second lane constraint 820 may be described by a third-degree polynomial. In addition to 3D modeling of first lane constraint 810 and/or second lane constraint 820, processing unit 110 may perform multi-frame estimation of host motion parameters, such as the speed, yaw and pitch rates, and acceleration of vehicle 200. Optionally, processing unit 110 may detect static and moving vehicles and their position, heading, speed, and acceleration, all relative to vehicle 200. Processing unit 110 may further determine a road elevation model to transform all of the information acquired from the plurality of images into 3D space.

[0155] Generally, as a default condition, vehicle 200 may travel relatively centered within first and second lane constraints 810 and/or 820. However, in some circumstances, environmental factors may make this undesirable or unfeasible, such as when objects or structures are present on one side of the road. Thus, there may be circumstances in which it may be desirable or advantageous for vehicle 200 to be closer to one lane constraint or the other. Processing unit 110 may also be configured to determine whether such circumstances, called lane offset conditions, exist based on the plurality of images.

[0156] To determine whether lane offset conditions exist, processing unit 110 may be configured to determine the presence of one or more objects 825 in the vicinity of vehicle 200. For example object 825 may comprise another vehicle, such as a car, truck or motorcycle traveling on roadway 800. Processing unit 110 may determine, for each object 825, an offset profile. The offset profile may include a determination of whether in its current and predicted position object 825 will be within a predefined range of vehicle 200. If processing unit 110 determines that object 825 is or will be within a predefined range of vehicle 200, processing unit 110 may determine whether there is enough space within first lane constraint 810 and second lane constraint 820, for vehicle 200 to bypass object 825. If there is not enough space, processing unit 110 may execute a lane change. If there is enough space, processing unit 110 may determine whether there is enough time to bypass object 825 before the distance between vehicle 200 and object 825 closes. If there is enough time, processing unit 110 may initiate and or activate the offset maneuver to bypass object 825. The offset maneuver may be executed by processing unit 110 so that the movement of vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 to 0.75 m/sec. The offset maneuver may be executed so that the maximum amplitude of the offset maneuver will occur when the gap be-

IPR2025-00943
Tesla EX1002 Page 1697

tween vehicle 200 and object 825 closes. If there is not enough time, processing unit 110 may initiate braking (e.g., by transmitting electronic signals to braking system 230) and then initiate offset maneuver.

[0157] After each offset profile is set for each object 825, processing unit may combine offset profiles such that the maximum lateral distance is kept from each object 825 that vehicle 200 is bypassing. That is, offset maneuvers may be modified and/or executed such that the lateral distance between vehicle 200 and each object 825 that vehicle 200 is bypassing is maximized.

[0158] Additionally or alternatively, object 825 may be a parked or stationary vehicle, a wall, or a person, such as pedestrian or cycling traffic. Processing unit 110 may determine a lane offset condition exists based on certain characteristics of object 825. For example, processing unit 110 may determine object 825 constitutes a lane offset condition if the height of object 825 exceeds a predetermined threshold, such as 10 cm from a road surface. In this manner, processing unit 110 may be configured to filter out objects that would not present a lane offset condition, such as small pieces of debris.

[0159] Additionally or alternatively, processing unit 110 may be configured to determine whether a lane offset condition exists on first vehicle side 802 based on a position of a target object 830 on vehicle first side 802. Target object 830 may include one or more objects that constitute a lane offset condition. Additionally or alternatively, target object 830 may not constitute a lane offset condition. For example, target object 830 may be a street sign or other object a sufficient distance from the travel lane or occurring with sufficient infrequency to not constitute a lane offset condition. Processing unit 110 may be configured to conclude a lane offset condition does not exist if such a target object 830 is detected.

[0160] Processing unit 110 may be configured to cause vehicle 200 to travel such that first distance 805 and second distance 810 are substantially different to avoid or stay farther away from object 825. For example, if a lane offset condition exists on first vehicle side 802, processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is greater than second distance 815. This is illustrated, for example, in FIG. 8A. Alternatively, if a lane offset condition exists on second vehicle side 812, processing unit 110 may create or execute an offset profile that is configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is less than second distance 815. This is illustrated, for example, in FIG 8B.

[0161] As another example, a parked vehicle may be located on one side of vehicle 200 and a pedestrian may be located on the other side of vehicle 200. In such a situation, system 100 may determine that is desirable to reduce the distance between vehicle 200 and the parked vehicle, in order to provide more space between vehicle 200 and the pedestrian. Accordingly, processing unit 110 may cause vehicle 200 to travel such that the distance between the parked vehicle is reduced.

[0162] The identity of first vehicle side 802 and second vehicle side 812 need not be static, such that at different points in time, processing unit 110 may identify first vehicle side 802 as the driver's side or as the passenger's side of vehicle 200. For example, the identity of first vehicle side 802 may change depending on the presence of one or more objects 825 on that side of vehicle 200, such that if one lane offset condition exists based on object 825, it exists on first vehicle side 802. In such circumstances, processing unit 110 may create or execute an offset profile that is configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805, which is greater than second distance 815, is provided toward object 825.

[0163] According to some embodiments, when a lane offset condition exists on first vehicle side 802, processing unit 110 may create an offset profile configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is at least 1.25 times greater than second distance 815. Additionally or alternatively, when a lane offset condition exists on first vehicle side 802, processing unit 110 may create an offset profile that is configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is at least 2 times greater than second distance 815.

[0164] In some embodiments, as shown in FIG. 8C, processing unit 110 may determine a lane offset condition exists on both first vehicle side 802 and on second vehicle side 812. Processing unit 110 may then create an offset profile for this offset condition that is configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is substantially the same as second distance 815.

[0165] Processing unit 110 may be configured so that the offset profiles for moving objects 825 as well as stationary objects 825 are prioritized. For example, processing unit 110 may prioritize offset profiles to keep a predefined distance between vehicle 200 and moving objects 825 and/or static objects 825. Processing unit 200 may prioritize these offset conditions over those offset conditions based on certain lane characteristics, such as a curve. The last priority lane offset condition may include mimicking a lead vehicle.

[0166] Processing unit 110 may factor offset profiles or offset conditions into a desired look-ahead point, which may be located at the middle of the lane defined by first lane constraint 810 and second lane constraint 820. Processing unit 110 may compute and implement a yaw rate command based on the look-ahead point and the offset profiles. The yaw rate command may be used to compute a torque command based on the yaw rate command and the yaw rate of vehicle 200 measured by using a gyroscope. Processing unit 110 may transmit the

IPR2025-00943
Tesla EX1002 Page 1698

torque command to a steering system (e.g., steering system 240) of vehicle 200.

[0167] According to some embodiments, input from a driver of vehicle 200, such as via user interface 170, may override processing unit 110. For example, a user input may override the control of processing unit 110 to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is greater than second distance 815.

[0168] Additionally or alternatively, processing unit 110 may be configured to determine, based on the plurality of images captured by image capture device 122, whether a lane offset condition is substantially non-existent. This may include circumstances in which one or more objects 825 below a certain size threshold are detected, one or more objects 825 detected are outside of a predetermined distance from the lane constraint, and/or when no objects 825 are detected. Under such circumstances, processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is equal to or substantially equal to second distance 815.

[0169] A lane offset condition may include a curved road, as illustrated in FIG. 8D. Processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is toward the outside of the curved road relative to vehicle 200 and second distance 815 is toward an inside of the curved road relative to vehicle 200. For example, first distance 805 may be less than second distance 815.

[0170] FIG. 9 is an exemplary block diagram of memory 140 and/or 150, which may stare instructions for performing one or more operations consistent with disclosed embodiments. As illustrated in FIG. 9, memory 140 may store one or more modules for performing the lane offset condition detection and responses described herein. For example, memory 140 may store a lane constraint module 910 and a lane offset module 920. Application processor 180 and/or image processor 190 may execute the instructions stored in any of modules 910 and 920 included in memory 140. One of skill in the art will understand that references in the following discussions to processing unit 110 may refer to application processor 180 and image processor 190 individually or collectively. Accordingly, steps of any of the following processes may be performed by one or more processing devices.

[0171] Lane constraint module 910 may store instructions which, when executed by processing unit 110, may detect and define first lane constraint 810 and second lane constraint 820. For example, lane offset module 910 may process the plurality of images received from at least one image capture device 122-124 to detect first lane constraint 810 and second lane constraint 820. As discussed above, this may include identifying painted lane lines and/or measuring a midpoint of a road surface.

[0172] Lane offset module 920 may store instructions which, when executed by processing unit 110, may detect the presence of a lane offset condition and/or identify whether lane offset condition is present on first vehicle side 802 and/or second vehicle side 812. For example, lane offset module 920 may process the plurality of images to detect the presence of object(s) 825 on first vehicle side 802 and/or second vehicle side 812. Lane offset module 920 may process the plurality of images to detect a curve in the road on which vehicle 200 is traveling. Additionally or alternatively, lane offset module 920 may receive information from another module or other system indicative of the presence of object(s) 825 and/or a curved lane. Lane offset module 920 may execute control to change the position of vehicle 200 to first lane constraint 810 and/or second lane constraint 820. For example, if lane offset module 920 determines a lane offset condition only on first vehicle side 802, lane offset module 920 may move vehicle 200 closer to second lane constraint 820. On the other hand, if lane offset module 920 determines a lane offset condition only on first vehicle side 812, lane offset module 920 may move vehicle 200 closer to first lane constraint 810.

[0173] Fig. 10 illustrates a process 1000 for navigating vehicle 200, consistent with disclosed embodiments. Process 1000 may identify lane constraints that define a lane of travel for vehicle travel. Process 1000 include causing vehicle 200 to travel closer to one or the other lane constraint, such as first lane constraint 810 or second lane constraint 820, in response to detecting a lane offset condition. In this manner, process 1000 may be used to drive vehicle 200 within lane constraints as well as account for environmental factors that may make it more desirable to shift vehicle 200 within first lane constraint 810 and second lane constraint 820.

[0174] At step 1010, process 1000 may include acquiring, using at least one image capture device 122, 124, and/or 126, a plurality of images of an area in the vicinity of vehicle 200. For example, processing unit 110 may receive the plurality of images may through data interface 128. For example, processing unit 110 may be configured to determine the presence of one or more objects 825 on the side of the road. For example, object 825 may be a parked or stationary vehicle, a wall, or a person, such as pedestrian or cycling traffic. Processing unit 110 may determine if a lane offset condition exists based on certain characteristics of object 825. For example, processing unit 110 may determine that object 825 constitutes a lane offset condition if the height of object 825 exceeds a predetermined threshold, such as 10 cm from a road surface. In this manner, processing unit 110 may be configured to filter out objects that would not present a lane offset condition, such as small pieces of debris.

[0175] Additionally or alternatively, processing unit 110 may be configured to determine whether a lane offset condition exists on first vehicle side 802 based on a position of a target object 830 on vehicle first side 802. Target object 830 may be one or more objects that constitute a lane offset condition. Additionally or alternatively, target object 830 may not constitute a lane offset condition. For

IPR2025-00943
Tesla EX1002 Page 1699

example, target object 830 may be a road barrier or a street sign. Processing unit 110 may be configured to conclude a lane offset condition does not exist if such a target object 830 is detected.

[0176] At step 1020, process 1000 may determine from the plurality of images a first lane constraint on first vehicle side 802 and a second lane constraint on second vehicle side 812. For example, processing unit 110 may be configured to determine first lane constraint 810 and second lane constraint 820 based on the plurality of images received via data interface 128.

[0177] At step 1030, process 1000 may determine whether a lane offset condition exists on first vehicle side 802. If a lane offset condition does exist on first vehicle side 802, at step 1040, process 1000 may include causing vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that second distance 815 is less than first distance 805. Processing unit 110 may be configured to cause vehicle 200 to travel such that first distance 805 and second distance 810 are substantially different to avoid or stay farther away from object 825. For example, if a lane offset condition exists on first vehicle side 802, processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is greater than second distance 815. This is illustrated, for example, in FIG. 8A. Alternatively, if a lane offset condition exists on second vehicle side 812, processing unit 110 may be configured to cause vehicle 200 to travel within first lane constraint 810 and second lane constraint 820 such that first distance 805 is less than second distance 815. This is illustrated, for example, in FIG 8B. For example, to navigate vehicle 200 according to process 1000, processing unit 110 may transmit electronic signals to one or more of throttling system 220, braking system 230, and/or steering system 240 of vehicle 200.

[0178] If a lane offset condition does not exist on first vehicle side 802, in step 1050, method may include determining whether lane offset condition exists on second vehicle side 812. If a lane offset condition exists on second vehicle side 812, the method may include causing vehicle 200 to travel between first lane constraint 810 and second lane constraint 820 such that first distance 805 is less than second distance 815.

Navigating a Vehicle to a Default Lane

[0179] System 100 may provide driver assist functionality that monitors the location of vehicle 200 in a current lane of travel and moves vehicle 200 into a predetermined default travel lane if the current lane of travel is not the predetermined default travel lane. The predetermined default travel lane may be defined by a user input and/or as a result of processing images received from at least one image capture device 122-126.

[0180] System 100 may also provide notification of an impending lane change of vehicle 200 to the predetermined default travel lane to the driver and passengers of vehicle 200 and/or other drivers. For example, system 100 may activate a turn signal positioned at or near one of bumper regions 210. Additionally or alternatively, system 100 may sound an audible notification, such as through speakers 360. Vehicle 200 may use system 100 to select a particular lane as predetermined default travel lane when vehicle is traveling. When the current lane within which vehicle 200 is traveling-e.g., the current lane of travel-differs from the predetermined default travel lane, system 100 may cause vehicle 200 to make a lane change, e.g., into the predetermined default travel lane.

[0181] FIG. 11 illustrates vehicle 200 traveling in a current lane of travel 1110 of roadway 800. According to some embodiments, vehicle 200 may have a predetermined default travel lane 1120. For example, predetermined default travel lane 1120 may be the right-most lane of roadway 800. Navigation system 100 may be configured to determine whether vehicle 200 is traveling in predetermined default travel lane 1120 and, if not, cause vehicle 200 to change to predetermined default travel lane 1120. For example, processing unit 110 may be configured to compare current lane of travel 1110 with predetermined default travel lane 1120 and cause vehicle 200 to return to predetermined default travel lane 1120.

[0182] According to some embodiments, processing unit 110 may detect an object 1130 in predetermined default travel lane 1120 based on a plurality of images acquired by one or more of image capture devices 122-126. Processing unit 110 may detect the position and speed of vehicle 200 relative to object 1130 by analyzing the images using techniques described in connection with FIGS. 5A-5D, above. According to some embodiments, processing unit 110 may determine whether vehicle 200 should pass object 1130. This may be desirable if, for example, the speed and/or acceleration of vehicle 200 relative to object 1130 exceeds a certain threshold.

[0183] FIG. 12A is an exemplary block diagram of memory 140, which may store instructions for performing one or more operations consistent with disclosed embodiments. As illustrated in FIG. 12A, memory 140 may store one or more modules for performing the default travel lane and current travel lane detection and responses described herein. For example, memory 140 may store a current lane detection module 1210, a default lane detection module 1220, and an action module 1230.

[0184] Current lane detection module 1210 may store instructions which, when executed by processing unit 110, may detect and define current lane of travel 1110. For example, processing unit 110 may execute current lane detection module 1212 to process the plurality of images received from at least one image capture device 122-124 and detect current lane of travel 1110. This may include detecting the relative position of vehicle 200 with respect to the edge of roadway 800, a midpoint of roadway 800, and/or lane markers, such as painted lane lines. Processing unit 110 may detect the position of vehicle

IPR2025-00943
Tesla EX1002 Page 1700

200 relative to roadway 800 by analyzing the images using techniques described in connection with FIGS. 5A-5D, above.

[0185] Default lane detection module 1220 may store instructions which, when executed by processing unit 110, may identify predetermined default travel lane 1120 and compare default lane detection module 1220 to current lane of travel 1110 identified by current lane detection module 1210. Default lane detection module 1220 may identify predetermined default travel lane 1120 based upon a user input and/or the plurality of images received from at least one image capture device 122-124. For example, a user may select a default lane through voice commands received by speakers 360 or by making a selection on a menu displayed on touch screen 320.

[0186] In some embodiments, predetermined default travel lane 1120 may be the right-most travel lane, or it may be the lane that has the more desirable features. For example, default lane detection module 1220 may include instructions to select the lane with less traffic, the lane in which vehicles are traveling closest to speed of vehicle 200, and/or a lane closest to an approaching exit or street to which vehicle 200 will turn. In some embodiments, predetermined default travel lane 1120 may be determined based on data received from position sensor 130, such as a global positioning system.

[0187] Action module 1230 may store instructions which, when executed by processing unit 110, may cause a response to the identification of current lane of travel 1110 and predetermined default travel lane 1120. For example, processing unit 110 may execute action module 1230 to cause vehicle 200 to change lanes if current lane of travel 1110 is not the same as predetermined default travel lane 1120. Additionally or alternatively, action module 1230 may initiate a notification of the lane change, such as activating turn signal 210 and or causing an audible announcement.

[0188] Action module 1230 may include instructions that cause processing unit 110 to determine whether vehicle 200 should move out of predetermined default travel lane 1120. For example, action module 1230 may determine that object 1130 in predetermined default travel lane 1120 should be bypassed by vehicle 200.

[0189] Action module 1230 may further include instructions that cause processing unit 110 to determine whether it is safe to make a lane change. For example, prior to executing a lane change, processing unit 110 may evaluate determine one or more characteristics of roadway 800 based on, for example, analysis of the plurality of images conducted using techniques described in connection with FIGS. 5A-5D, above. If the analysis indicates conditions are not safe to change lanes, processing unit 110 may determine that vehicle 200 should stay in the current lane.

[0190] Action module 1230 may include additional instructions for manipulating vehicle 200 to maneuver out of and return to predetermined default travel lane 1120, such as for bypassing object 1130. When executing these instructions, processing unit 110 may estimate a location of predetermined default travel lane 1120 using estimated motion of vehicle 200. This may be advantageous in circumstances in which predetermined default travel lane 1120 does not remain detectable by image capture devices 122-124.

[0191] Processing unit 110 may be configured to execute an offset profile that offsets vehicle 200 from predetermined default travel lane 1120 by a full lane width, which may be estimated based on the distance between detected lane markers. The offset maneuver defined by offset profile may be executed by processing unit 110 so that the movement of vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 to 0.75 m/sec. During the offset maneuver, processing unit 110 may execute instructions from action module 1230 that switch to following the center of current lane of travel 1110. These instructions may include gradually transitioning from following the predicted lane with offset to following current lane of travel 1110. A gradual transition may be accomplished by defining desired look-ahead point so that it is not allowed to change in time more than the recommended slew rate, which may be around 0.25 m/sec.

[0192] The instructions for transitioning vehicle 200 from predetermined default travel lane 1120 may be used to transition vehicle 200 back into predetermined default travel lane. For example, processing unit 110 may be configured to execute an offset profile that offsets vehicle 200 from current lane of travel 1110 by a full lane width, which may be estimated based on the distance between detected lane markers. The offset maneuver defined by offset profile may be executed by processing unit 110 so that the movement of vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 to 0.75 m/sec. During the offset maneuver, processing unit 110 may execute instructions from action module 1230 that switch to following the center of predetermined default travel lane 1120. These instructions may include gradually transitioning from following the predicted lane with offset to following predetermined default travel lane 1120. A gradual transition may be accomplished by defining desired look-ahead point so that it is not allowed to change in time more than the recommended slew rate, which may be around 0.25 m/sec.

[0193] At one or more of the steps for transitioning vehicle 200 between two lanes, processing unit 110 may compute and implement a yaw rate command based on the look-ahead point and the offset profiles. The yaw rate command may be used to compute a torque command based on the yaw rate command and the yaw rate of vehicle 200 measured using a gyroscope. Processing unit 110 may transmit a torque command to a steering system (e.g., steering system 240) of vehicle 200.

[0194] FIG. 12B illustrates a process 1250 for navigating vehicle 200 to a default lane consistent with disclosed embodiments. According to some embodiments, proc-

IPR2025-00943
Tesla EX1002 Page 1701

ess 1250 may be implemented by one or more components of navigation system 100, such as at least one processing unit 110. Process 1250 may identify current lane of travel 1110 of vehicle 200 and compare it with predetermined default travel lane 1120 to determine whether vehicle 200 is traveling within predetermined default travel lane 1120. If vehicle 200 is not traveling in predetermined default travel lane 1120, processing unit 110 may cause vehicle 200 to change lanes so that it navigates to default travel lane 1120.

[0195] At step 1260, one or more of image capture devices 122-124 may acquire a plurality of images of an area in a vicinity of vehicle 200. Processing unit 110 may receive the plurality of images via data interface 128.

[0196] At step 1270, processing unit 110 may execute current lane detection module 1210 detect and define current lane of travel 1110. For example, processing unit 110 may process the plurality of images to, for example, detect the relative position of vehicle 200 with respect to the edge of roadway 100, a midpoint of roadway 100, and/or painted lane markers. Processing unit 110 may detect the position of vehicle 200 relative to roadway 800 by analyzing the images using techniques described in connection with FIGS. 5A-5D, above.

[0197] At step 1280, processing unit 110 may execute default lane detection module 1220 to identify predetermined default travel lane 1120 and compare default lane detection module 1220 to current lane of travel 1110 identified by current lane detection module 1210. Default lane detection module 1220 may identify predetermined default travel lane 1120 based upon a user input and/or the plurality of images received from at least one image capture device 122-124.

[0198] For example, as discussed above, processing unit 110 may determine predetermined default travel lane 1120 based on any number of factors. Predetermined default travel lane 1120 may be the right-most lane among the plurality of travel lanes 1110 and 1120. Additionally or alternatively, predetermined default travel lane 1120 may be determined based on an input received from a user via user interface 170, by processing the plurality of images, and/or based on data received from position sensor 130, such as a global positioning system. Predetermined default travel lane 1120 may be determined once, once per vehicle trip, or on regular intervals. For example, predetermined default travel lane 1120 may be determined dynamically in response to conditions at a given location.

[0199] If current lane 1110 is not the same as predetermined default travel lane 1120, processing unit 110 may execute action module 1230 to cause vehicle 100 to navigate to default travel lane 1120. For example, processing unit 110 may transmit electronic signals to one or more of throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to trigger the desired response. For example, processing unit 110 may cause steering system 240 to turn the steering wheel of vehicle 200 to achieve a rotation of a predetermined angle.

[0200] Further, prior to executing a lane change, processing unit 110 may determine whether it is safe (e.g., there are no vehicles or objects in the way) to change lanes. As discussed above, processing unit 110 may evaluate one or more characteristics of roadway 800 based on, for example, analysis of the plurality of images conducted using techniques described in connection with FIGS. 5A-5D, above. If conditions are not safe to change lanes, processing unit 110 may determine that vehicle 200 should stay in the current lane.

[0201] In some embodiments, execution of action module 1230 may initiate a notification of the lane change. For example, processing unit 110 may cause a signal to activate a turn signal and/or cause an audible announcement via speakers 360. Accordingly, such notification may occur prior to vehicle 200 changing lanes. In other embodiments, execution of action module 1230 may initiate a notification to a driver of vehicle 200 that vehicle 200 is not traveling in a default lane without causing the lane change to occur. Accordingly, in such an embodiment, the driver of vehicle 200 may determine whether to manually take action and steer vehicle 200 to change lanes.

[0202] In some embodiments, system 100 may notify the user of vehicle 200 and/or other drivers in the vicinity of vehicle 200 that vehicle 200 will be changing lanes if vehicle 200 will be moved to predetermined default travel lane 1120. For example, processing unit 110 may activate a turn signal of vehicle 200 prior to causing vehicle 200 to change lanes. Additionally or alternatively, an audible announcement to notify other drivers in the vicinity may be caused by processing unit 110. The audible announcement may be transmitted via, for example, Bluetooth, to receivers included in other nearby vehicles. Additionally or alternatively, at least one processing unit 110 may cause an audible indicator to notify the driver of vehicle 200 prior to causing vehicle 200 to switch lanes. In this manner, driver of vehicle 200 can anticipate the control system 100 will exert on vehicle 200.

Controlling Velocity of a Turning Vehicle

[0203] System 100 may provide driver assist functionality that controls the velocity (i.e., speed and/or direction) of vehicle 200 in different scenarios, such as while making a turn (e.g., through a curve). For example, system 100 may be configured to use a combination of map, position, and/or visual data to control the velocity of vehicle 200 when turning. In particular, system 100 may control the velocity of vehicle 200 as it approaches a curve, while navigating a curve, and/or as it exits a curve. System 100 may consider different factors to control the velocity depending on the position of vehicle 200 relative to the curve at a given time. For example, system 100 may initially cause vehicle 200 to slow down in response to learning that a curve is located a certain distance ahead of vehicle 200. While vehicle 200 turns through

IPR2025-00943
Tesla EX1002 Page 1702

the curve, however, system 100 may cause vehicle 200 to accelerate and/or charge direction based on an analysis of the characteristics of the curve.

[0204] FIGS. 13A and 13B are diagrammatic representations of a vehicle (e.g., vehicle 200) approaching, and navigating, a curve with one or more characteristics on road 1300, consistent with the disclosed embodiments. As illustrated in FIG. 13A, visual indicators, such as road sign 1302 located near road 1300, may inform vehicle 200 that a curve lies ahead. In response to the warning provided by road sign 1302, a driver of vehicle 200 and/or a system 100 providing driver assist functionality of vehicle 200 may adjust the velocity of vehicle 200 (e.g., by slowing down and/or steering) to safely navigate the curve.

[0205] As vehicle 200 approaches the curve, system 100 may detect the lane width 1306 and the curve radius 1308 by using image capture devices 122 and/or 124 (which may be cameras) of vehicle 200. Based on the curve radius 1308, the lane width 1306, and other characteristics associated with the curve detected using image capture devices 122 and/or 124, system 100 may adjust the velocity of vehicle 200. Specifically, system 100 may adjust the velocity as vehicle 200 approaches the curve and/or while vehicle 200 navigates the curve. System 100 may also adjust the velocity as vehicle 200 exits the curve, in response to, for example, traffic light 1304 appearing after the curve. Further detail regarding adjustments to the velocity of vehicle 200 and detection of the lane width 1306, the radius of curvature or curve radius 1308, and characteristics associated with the curve are described in connection with FIGS. 14 and 15, below.

[0206] In some embodiments, system 100 may recognize a curve to be navigated based on map data and/or vehicle position information (e.g., GPS data). System 100 may determine an initial target velocity for the vehicle based on one or more characteristics of the curve as reflected in the map data. System 100 may adjust a velocity of the vehicle to the initial target velocity and determine, based on one or more images acquired by one or more of image capture devices 122-126, one or more observed characteristics of the curve. System 100 may determine an updated target velocity based on the one or more observed characteristics of the curve and adjust the velocity of the vehicle to the updated target velocity.

[0207] FIG. 14 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments. As indicated in FIG. 14, memory 140 may store a curve recognition module 1402, a characteristic observation module 1404, and a velocity module 1406. The disclosed embodiments are not limited to any particular configuration of memory 140. Further, processing unit 110 may execute the instructions stored in any of modules 1402-1406 included in memory 140.

[0208] In one embodiment, curve recognition module 1402 may store software instructions (such as computer vision software) which, when executed by processing unit 110, detects a curve by, for example, analyzing one or more sets of images and/or using map and/or position data (e.g., data stored in map database 160). The image sets may be acquired by one or more of image capture devices 122, 124, and 126. Based on the analysis, system 100 (e.g., via processing unit 110) may cause a change in velocity of vehicle 200. For example, processing unit 110 may cause vehicle 200 to reduce its speed by a predetermined amount and/or adjust its direction by a particular angle if processing unit 110 detects the presence of a curve located within a minimum threshold distance ahead of vehicle 200.

[0209] In one embodiment, characteristic observation module 1404 may store software instructions (such as computer vision software) which, when executed by processing unit 110, observes characteristics associated with a curve detected by system 100 (e.g., via execution of curve recognition module 1402). Processing unit 110 may observe characteristics associated with a detected curve by analyzing one or more sets of images acquired by one or more of image capture devices 122, 124, and 126 and/or using map and/or position data (e.g., data stored in map database 160). The analysis may yield information associated with the curve, such as curve radius 1308, lane width 1306, a degree of curvature, a rate of change in curvature, a degree of banking, a length or arc length of the curve, and the like. These characteristics may include estimated values (e.g., curve characteristics calculated using preexisting data and/or mathematical models based on such data) and actual values (e.g., curve characteristics calculated by analyzing captured images of the curve). Based on the analysis, processing unit 110 may cause a change in velocity of vehicle 200. For example, processing unit 110 may cause vehicle 200 to reduce its speed and/or adjust its direction by a particular angle in view of a high rate of change in curvature associated with a curve located 25 meters ahead of vehicle 200.

[0210] In one embodiment, velocity module 1406 may store software instructions configured to analyze data from one or more computing and electromechanical devices configured to determine a target velocity and cause a change in velocity of vehicle 200 to the target velocity. For example, processing unit 110 may execute velocity module 1406 to calculate a target velocity for vehicle 200 based on data derived from execution of curve recognition module 1402 and characteristic observation module 1404. Such data may include, for example, an initial target position and initial velocity, an updated target position and updated velocity, a final target position and final velocity, and the like. In addition, processing unit 110 may calculate a velocity for vehicle 200 based on input from other systems of vehicle 200, such as a throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. Based on the calculated target velocity, processing unit 110 may transmit electronic signals (e.g., via a controller area network bus ("CAN bus")) to throttling

IPR2025-00943
Tesla EX1002 Page 1703

system 220, braking system 230, and/or steering system 240 of vehicle 200 to trigger a change in the speed and/or direction of vehicle 200 by, for example, physically depressing the brake, easing up off the accelerator of vehicle 200, or steering vehicle 200 in a particular direction (e.g., at a particular angle). An increase in speed may be associated with a corresponding acceleration (e.g., between 0.2 m/sec^2 and 3.0 m/sec^2). Conversely, a decrease in speed may be associated with a corresponding deceleration. In addition, the acceleration or deceleration may be based on the road type (e.g., highway, city street, country road, etc.), the presence of any speed constraints nearby (e.g., sharp curves, traffic lights), as well as a road width or lane width 1306. For purposes of this disclosure, deceleration may refer to an acceleration having a negative value.

[0211] FIG. 15 is a flow chart showing an exemplary process 1500 for controlling the velocity of a vehicle based on a detected curve and observed characteristics of the curve and/or map data regarding the curve consistent with disclosed embodiments. At step 1510, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120. For instance, a camera included in image acquisition unit 120 (such as image capture device 122) may capture a plurality of images and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive more than one plurality of images via a plurality of data interfaces. For example, processing unit 110 may receive a plurality of images from each of image capture devices 122, 124, 126, each of which may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0212] At step 1520, processing unit 110 may execute curve recognition module 1502 to detect a curve located ahead of vehicle 200 based on, for example, map and/or position data (e.g., data stored in map database 160). Such data may indicate the location of a curve that is present on a given road. The data may indicate the absolute position of a particular curve (e.g., via GPS coordinates) or the relative position of the curve (e.g., by describing the curve relative to vehicle 200, road sign 1302, traffic light 1304, and/or other descriptive landmarks). In some embodiments, processing unit 110 may execute curve recognition module 1502 to detect the curve by analyzing the plurality of images. The analysis may be conducted using techniques described in connection with FIGS. 5A-5D and 6, above.

[0213] At step 1530, processing unit 110 may execute velocity module 1406 to determine an initial target velocity for vehicle 200 based on, for example, data derived from execution of curve recognition module 1402. The initial target velocity may reflect a speed and a direction for vehicle 200 to safely enter an upcoming curve and may be determined using techniques described in connection with FIG. 14, above. For example, velocity module 1406 may determine the initial target velocity based on known or estimated characteristics associated with the upcoming curve, such as the curve radius 1308 and the degree of banking. Velocity module 1406 may calculate the initial target velocity in accordance with mathematical models and/or formulae and based on various constraints. In some embodiments, for example, velocity module 1406 may determine the initial target velocity based on the presence of traffic lights (such as traffic light 1304) before, during, or after a curve, weather conditions (e.g., light rain, heavy rain, windy, clear, etc.), road conditions (e.g., paved road, unpaved road, potholes, etc.), posted speed limits, the presence of vehicles or pedestrians nearby, etc. The initial target velocity may also be based on a lateral acceleration limit associated with vehicle 200. In some embodiments, the lateral acceleration limit may be in the range of 0.2 m/sec^2 and 3.0 m/sec^2 and may be adjusted depending on the curve characteristics, the road width, or lane width 1306. Additionally, the initial target velocity may be based on a focal range associated with one or more of cameras, such as image capture devices 122 and 124.

[0214] For example, velocity module 1406 may determine the initial target velocity based on the radius of curvature of the road ahead of vehicle 200. Velocity module 1406 may determine the radius of curvature at various distances from vehicle 200 and store the information in a vector R (1:n). In one embodiment, velocity module 1406 may determine the radius of curvature at: (1) the current position of vehicle 200, (2) a distance D_{end} from vehicle 200, where D_{end} represents the distance ahead of vehicle 200 at which knowledge of the road characteristics allows system 100 to control navigational responses in vehicle 200 such that the driver may feel comfortable, and (3) at incremental distances from vehicle 200, such as every 1 meter, etc. The value of distance D_{end} may depend on the current speed of vehicle 200, V_{ego}. For example, the greater the value of V_{ego} (e.g., the faster vehicle 200 is traveling), the greater the value of D_{end} (e.g., the longer the distance may be taken into account to provide system 100 with sufficient time to provide a navigational response, such as comfortably reducing the speed of vehicle 200 based on an approaching sharp turn or curve). In one embodiment, D_{end} may be calculated according to the following equation: $D_{end} = (V_{ego}^2/2)/(0.7/40*V_{ego}+0.5)$. Thus, the vector $D(1{:}n) = [1{:}1{:}n]$, where $n = \text{floor}(D_{end})$.

[0215] In some embodiments, velocity module 1406 may determine the initial target velocity based on a lateral acceleration constraint, a_{max}, e.g., the maximum allowed lateral acceleration, a_{max} may fall in the range of 1.5 to 3.5 meters / sec^2. Based on a_{max} and the radius of curvature vector R (discussed above), velocity module 1406 may limit the speed of vehicle 200 (e.g., at different points or distances from the current position of vehicle 200) to a speed v_{lim}, calculated as: $v_{lim}(1{:}n)=\text{sqrt}(a_{max}*R(1{:}n))$. In addition, a deceleration constraint that follows from

IPR2025-00943
Tesla EX1002 Page 1704

the determination of v_{lim} may be calculated as: dec_{lim} = min{{$(v_{lim}(1:n)^2 - V_{ego}^2)/2$), max{$d_{min}$,$D(1:n)-2*V_{ego}$}}, where d_{min} is equal to 40 meters. The value of d_{min} may represent a point at which deceleration feels comfortable for the driver of vehicle 200.

[0216] In some embodiments, velocity module 1406 may determine the initial target velocity based on a curvature slew rate constraint, $slew_{max}$, which may be based on the curvature rate of change. $slew_{max}$ may fall in the range of 0.5 to 1.5 meters / sec^3. Velocity module 1406 may determine v_{lim} according the following equation: $v_{lim}(1:n) = (slew_{max}/\text{abs}\{\text{diff}\{1/R(1:n)\}\})^{1/3}$. Further, velocity module 1406 may determine dec_{lim} according to the equation: dec_{lim} = min{{$(v_{lim}(1:n)^2 - V_{ego}^2)/2$), max {$d_{min}$, $D(1:n) - 2*V_{ego}$}}, where d_{min} is equal to 40 meters.

[0217] In some embodiments, velocity module 1406 may determine a comfort constraint such that the driver of vehicle 200 is comfortable while vehicle 200 is undergoing various navigational responses, such as making a turn or navigating a curve. Velocity module 1406 may determine v_{lim} according to the following equation: $v_{lim}(1:n) = 4/500*R(1:n) + 22.5$. Further, velocity module 1406 may determine dec_{lim} according to the equation: dec_{lim} = min{{$(v_{lim}(1:n)^2 - V_{ego}^2)/2$), max {$d_{min}$, $D(1:n) - 2*V_{ego}$}}, where d_{min} is equal to 40 meters.

[0218] Velocity module 1406 may merge the three constraints determined above-the lateral acceleration constraint, the curvature slew rate constraint, and comfort constraint-by using the smallest value of dec_{lim} calculated in association with each constraint. After merging the constraints, velocity module 1406 may arrive at the initial target velocity by determining how much to adjust the speed of vehicle 200, denoted v_{com}, according to the following equation: $v_{com} = v_{com(prev)} + dec_{lim}*\Delta t$, where $v_{com(prev)}$ represents the prior speed of vehicle 200 (e.g., a previously determined speed, the speed of vehicle 200 prior to any adjustment by velocity module 1406, etc.), and Δt represents the time between the determination of v_{com} and $v_{com(prev)}$.

[0219] At step 1540, processing unit 110 may execute velocity module 1406 to adjust a velocity of vehicle 200 to the initial target velocity determined at step 1530. Based on the current velocity of vehicle 200 and the initial target velocity determined at step 1530, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to change the velocity of vehicle 200 to the initial target velocity. Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240. For example, processing unit 110 may transmit electronic signals that cause system 100 to physically depress the brake by a predetermined amount or ease partially off the accelerator of vehicle 200. Further, processing unit 110 may transmit electronic signals that cause system 100 to steer vehicle 200 in a particular direction. In some embodiments, the magnitude of acceleration (or de-acceleration) may be based on the difference between the current velocity of vehicle 200 and the initial target velocity. Where the difference is large (e.g., the current velocity is 10 km/hr greater than the initial target velocity), for example, system 100 may change the current velocity by braking in a manner that maximizes deceleration and results in vehicle 200 achieving the initial target velocity in the shortest amount of time safely possible. Alternatively, system 100 may apply the brakes in a manner that minimizes deceleration and results in vehicle 200 achieving the initial target velocity gradually. More generally, system 100 may change the current velocity to the initial target velocity according to any particular braking profile (e.g., a braking profile calling for a high level of braking for the first two seconds, and a low level of braking for the subsequent three seconds). The disclosed embodiments are not limited to any particular braking profile or manner of braking.

[0220] At step 1550, processing unit 110 may execute characteristic observation module 1404 to determine one or more characteristics of a curve located ahead of vehicle 200 based on, for example, analysis of the plurality of images conducted using techniques described in connection with FIGS. 5A-5D, above. The characteristics may include, for example, estimated and/or actual values associated with curve radius 1308, a road width or lane width 1306, a degree of curvature, a rate of change in curvature, a degree of banking, and the like. The characteristics may also include, for example, information conveyed by one or more road signs (such as road sign 1302). Furthermore, the characteristics may include the length or arc length of the detected curve, which may be a function of the speed of vehicle 200 as it approaches the curve. Processing unit 110 may execute characteristic observation module 1404 to calculate values associated with these characteristics (e.g., curve radius 1308, the degree of curvature, the rate of change in curvature, the arc length of the curve, etc.) based on mathematical models and/or formulae. In some embodiments, processing unit 110 may execute characteristic observation module 1404 to construct a mathematical model expressing the curvature of the curve as a function of distance (e.g., distance traveled). Such a model may be subject to constraints such as a maximum lateral acceleration constraint (e.g., in the range of 1.5 to 3.0 m/sec^2) and/or a maximum lateral acceleration derivative constraint (e.g., in the range of 0.8 to 1.2 m/sec^2). These constraints may yield a maximum speed for vehicle 200 as a function of distance traveled.

[0221] At step 1560, processing unit 110 may execute velocity module 1406 to calculate an updated target velocity based on the characteristics determined at step 1550. The updated target velocity may reflect an updated speed and/or direction for vehicle 200 to safely perform any combination of: (i) entering an upcoming curve, (ii) navigating through a curve, or (iii) exiting a curve. The updated target velocity may be determined using techniques described in connection with FIG. 14 and step 1530, above. For example, velocity module 1406 may determine the updated target velocity based on charac-

IPR2025-00943
Tesla EX1002 Page 1705

teristics associated with a curve and mathematical models and/or formulae.

[0222] At step 1570, processing unit 110 may execute velocity module 1406 to adjust the velocity of vehicle 200 to the updated target velocity determined at step 1560. This adjustment may be accomplished using techniques described in connection with step 1540, above. For example, based on the current velocity of vehicle 200 and the updated target velocity determined at step 1560, processing unit 110 may transmit electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to change the velocity of vehicle 200 to the updated target velocity. At this point, the current velocity of vehicle 200 is not necessarily the same as the initial target velocity determined at step 1530. For example, processing unit 110 may have determined the updated target velocity at step 1560 while in the process of adjusting the velocity of vehicle 200 to the initial target velocity. This may be true in the case where the difference in the velocity of vehicle 200 and the initial target velocity is large (e.g., greater than 10 km/hr) and the associated acceleration is small.

[0223] In some embodiments, system 100 may make regular updates to the target velocity based on continued observations of characteristics of a curve (e.g., based on image data depicting portions of the curve that vehicle 200 is approaching). Additionally, after system 100 has navigated vehicle 200 through a curve, system 100 may cause vehicle 200 to accelerate to a new target velocity suitable for traveling a segment of the road that does not have curves, as determined by analyzing map data, positional data, and/or image data acquired by system 100.

Mimicking A Leading Vehicle

[0224] System 100 may provide driver assist functionality that causes vehicle 200 to mimic (or decline to mimic) a leading vehicle in different scenarios, such as when the leading vehicle switches lanes, accelerates, or makes a turn. For example, system 100 may be configured to detect the leading vehicle by analyzing a plurality of images and determine one or more actions taken by the leading vehicle. In some scenarios, such as when the leading vehicle is turning at an intersection, system 100 may be configured to cause vehicle 200 to decline to mimic the turn. In other scenarios, such as when the leading vehicle is turning at an intersection without changing lanes within the intersection, system 100 may be configured to cause vehicle 200 to mimic the turn. System 100 may also be configured to cause vehicle 200 to mimic one or more actions of the leading vehicle based on a navigation history of the leading vehicle.

[0225] FIGS. 16A and 16B are diagrammatic representations of a primary vehicle 200a mimicking one or more actions of a leading vehicle 200b on road 1600a, consistent with the disclosed embodiments. As illustrated in FIGS. 16A and 16B, primary vehicle 200a may be trailing leading vehicle 200b while traveling in the same lane of road 1600a. The lane may have a left edge 1620, a right edge 1630, and a midpoint 1610, with w_1 and w_2 indicating the left and right halves of the lane, respectively. Primary vehicle 200a and leading vehicle 200b may be positioned in the lane such that neither vehicle is centered on the midpoint 1610. As shown in FIG. 16A, for example, primary vehicle 200a may be centered to the right of midpoint 1610 such that the distance c_1 between primary vehicle 200a and left edge 1620 is greater than the distance c_2 between primary vehicle 200a and right edge 1630. In contrast, leading vehicle 200b may be centered to the left of midpoint 1610 such that the distance c_3 between leading vehicle 200b and left edge 1620 is less than the distance c_4 between leading vehicle 200b and right edge 1630. Different configurations of the positions of primary vehicle 200a and leading vehicle 200b are possible. For example, primary vehicle 200a and leading vehicle 200b may be positioned anywhere on road 1600a. Further, although FIG. 16A depicts leading vehicle 200b as having a camera system, the disclosed embodiments are not limited to a configuration in which leading vehicle 200b includes a camera system.

[0226] System 100 of primary vehicle 200a may cause primary vehicle 200a to mimic one or more actions of leading vehicle 200b. As shown in FIG. 16B, for example, primary vehicle 200a may mimic a lane shift performed by leading vehicle 200b. After mimicking the lane shift, distance c_1 may be equal to distance c_3, and distance c_2 may be equal to distance c_4. In some embodiments, primary vehicle 200a may mimic a lane shift performed by leading vehicle 200b in a way such that distances c_1 and c_3, and c_2 and c_4, are unequal. For example, in the case where primary vehicle 200a mimics a leftward lane shift of primary vehicle 200b and is positioned (after the lane shift) to the left of primary vehicle 200a, c_1 will be less than c_3 and c_2 will be greater than c_4. Alternatively, where primary vehicle 200a mimics a leftward lane shift of primary vehicle 200b and is positioned to the right of primary vehicle 200a, c_1 will be greater than c_3 and c_2 will be less than c_4.

[0227] FIGS. 16C and 16D are diagrammatic representations of a primary vehicle 200a mimicking one or more actions of a leading vehicle 200b on road 1600b, consistent with the disclosed embodiments. As illustrated in FIG. 16C, primary vehicle 200a may be traveling behind leading vehicle 200b, which may make a left turn at an intersection. System 100 of primary vehicle 200a may cause primary vehicle 200a to mimic the left turn of leading vehicle 200b, as shown in FIG. 16D. System 100 may cause primary vehicle 200a to mimic the left turn based on various factors, such as whether leading vehicle 200b changes lanes while turning within the intersection, the navigation history associated with leading vehicle 200b, the relative difference in speed between primary vehicle 200a and leading vehicle 200b, and the like. Additionally, primary vehicle 200a may mimic leading vehicle 200b under various conditions, such as based on whether leading vehicle 200b is traveling along the same route or

IPR2025-00943
Tesla EX1002 Page 1706

the same portion of a route that primary vehicle 200a is traveling along. Further detail regarding scenarios in which system 100 may cause primary vehicle 200a to mimic one or more actions of leading vehicle 200b are described in connection with FIGS. 17-22, below.

[0228] In some embodiments, system 100 may be configured to determine whether to mimic a particular movement of a leading vehicle. For example, system 100 may determine that certain movements of the leading vehicle do not need to be mimicked because the movements do not affect the

[0229] course of the leading vehicle. For example, system 100 may determine that small changes in movements, such as moving within a particular lane, do not need to be mimicked. In other embodiments, system 100 may implement a smoothing operation to filter out small movements of a leading vehicle. The smoothing operation may cause the primary vehicle to implement more significant and/or important movements (e.g., changing lanes) of a leading vehicle while filtering out smaller movements not affecting the overall course of the leading vehicle.

[0230] FIG. 17 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations consistent with the disclosed embodiments. As shown in FIG. 17, memory 140 may store a lane constraint module 1702, an action detection module 1704, and an action response module 1706. The disclosed embodiments are not limited to any particular configuration of memory 140. Further, processing unit 110 may execute the instructions stored in any of modules 1702-1706 included in memory 140.

[0231] In one embodiment, lane constraint module 1702 may store software instructions (such as computer vision software) which, when executed by processing unit 110, determine one or more constraints associated with a lane that primary vehicle 200a is traveling in. The lane constraints may include a midpoint 1610, a left edge 1620, and a right edge 1630 of the lane. The lane constraints may also include one or more lines or other symbols marked on the surface of road 1600a or 1600b. Processing unit 110 may execute lane constraint module 1702 to determine the constraints by, for example, analyzing one or more sets of images and/or using map and/or position data (e.g., data stored in map database 160). The image sets may be acquired by one or more of image capture devices 122, 124, and 126. Based on the analysis, system 100 (e.g., via processing unit 110) may cause primary vehicle 200a to travel within a lane defined by left edge 1620 and right edge 1630. In some embodiments, system 100 may cause primary vehicle 200a to move away from left edge 1620 or right edge 1630 if one of distance c_1 or c_2 is less than a predetermined threshold.

[0232] In one embodiment, action detection module 1704 may store software instructions (such as computer vision software) which, when executed by processing unit 110, detects one or more actions taken by leading vehicle 200b traveling in front of primary vehicle 200a. Leading vehicle 200b may be traveling in the same or different lane as primary vehicle 200a. Processing unit 110 may execute action detection module 1704 to detect one or more actions taken by leading vehicle 200b by, for example, analyzing one or more sets of images acquired by one or more of image capture devices 122, 124, and 126. Based on the analysis, processing unit 110 may determine, for example, that leading vehicle 200b has shifted lanes, made a turn, accelerated, decelerated, applied its brakes, and/or the like. As another example, leading vehicle 200b may perform a maneuver such that a first offset distance on a side of leading vehicle 200b adjacent to a first lane constraint is different from a second offset distance on a side of leading vehicle 200b adjacent to a second lane constraint. For example, processing unit 110 may perform the analysis based on the techniques described in connection with FIGS. 5A-5D, above.

[0233] In one embodiment, action response module 1706 may store software instructions (such as computer vision software) which, when executed by processing unit 110, determines one or more actions for primary vehicle 200a to take based on one or more actions taken by leading vehicle 200b detected by system 100 (e.g., via execution of action detection module 1704). For example, processing unit 110 may execute action response module 1706 to determine whether to mimic one or more actions of leading vehicle 200b. This determination may be based on the nature of leading vehicle 200b's actions (e.g., turn, lane shift, etc.), information associated with primary vehicle 200a (e.g., speed, distance c_1 and c_2 from the lane edges, etc.), road and environmental conditions (e.g., potholes, heavy rain or wind, etc.), and the like. In the case where processing unit 110 determines to mimic one or more actions of leading vehicle 200b, processing unit 110 may accomplish this by, for example, transmitting electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to trigger a turn, lane shift, change in speed, and/or change in direction of primary vehicle 200a. Processing unit 110 may use the techniques described in connection with FIGS. 4-7, above, to cause primary vehicle 200a to mimic the one or more actions. To mimic the one or more actions, system 100 may provide control signals to one or more of throttling system 220, braking system 230, and steering system 240 to navigate vehicle 200 (e.g., by causing an acceleration, a turn, a lane shift, etc.). Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240.

[0234] FIG. 18 is a flow chart showing an exemplary process for causing primary vehicle 200a to mimic one or more actions of leading vehicle 200b consistent with disclosed embodiments. At step 1810, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120. For instance, a camera included in image ac-

IPR2025-00943
Tesla EX1002 Page 1707

quisition unit 120 (such as image capture device 122) may capture a plurality of images and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive more than one plurality of images via a plurality of data interfaces. For example, processing unit 110 may receive a plurality of images from each of image capture devices 122, 124, 126, each of which may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0235] At steps 1820 and 1830, processing unit 110 may execute lane constraint module 1702 to determine constraints for each side of primary vehicle 200a associated with a lane that primary vehicle 200a is traveling in. Processing unit 110 may determine the constraints by analyzing the images received at step 1810. Thus, at step 1820, processing unit 110 may determine a first lane constraint on a first side of primary vehicle 200a (e.g., the distance between the left side of primary vehicle 200a and left edge 1620 of the lane, c_1) and, at step 1830, a second lane constraint on a second side opposite the first side (e.g., the distance between the right side of primary vehicle 200a and right edge 1620 of the lane, c_2). In some embodiments, processing unit 110 may determine the constraints by using map and/or position data (e.g., data stored in map database 160) indicating the position of primary vehicle 200a relative to midpoint 1610, left edge 1620, and right edge 1630 of the lane. Processing unit 110 may determine the constraints based on both an analysis of the images received at step 1810 and map and/or position data; doing so may increase the confidence level associated with the constraints.

[0236] At step 1840, processing unit 110 may execute lane constraint module 1702 to cause primary vehicle 200a to travel within the lane constraints determined at steps 1820 and 1830. For example, the left side of primary vehicle 200a may be positioned to the left of left edge 1620, or the right side of primary vehicle 200a may be positioned to the right of right edge 1630. In such a scenario, processing unit 110 may transmit electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to cause primary vehicle 200a to adjust its position and travel within the lane constraints.

[0237] At step 1850, processing unit 110 may execute action detection module 1704 to locate a leading vehicle 200b within the images received at step 1810 by analyzing the images. Processing unit 110 may detect the presence of vehicles in the images using the techniques described in connection with FIGS. 5A and 5B, above. For each detected vehicle, processing unit 110 may construct a multi-frame model of the detected vehicle's position, velocity (e.g., speed and direction), and acceleration relative to primary vehicle 200a. Processing unit 110 may use the techniques described in connection with FIGS. 5A and 5B, above, to construct such a model. In addition, processing unit 110 may determine a lane of travel associated with each detected vehicle, based on, for example, lane or other road markings and lane constraints determined at steps 1820 and 1830. Processing unit 110 may determine leading vehicle 200b to be the detected vehicle that is closest to, and is traveling in the same lane as, primary vehicle 200a. At step 1860, processing unit 110 may execute action detection module 1704 to determine one or more actions taken by the leading vehicle 200b detected within the images at step 1850. For example, processing unit 110 may determine that leading vehicle 200b has shifted lanes, made a turn, accelerated, decelerated, and/or the like. In addition, processing unit 110 may analyze the images to determine characteristics associated with leading vehicle 200b, such as its speed and position (e.g., relative to primary vehicle 200a). These determinations may be based on the techniques described in connection with FIGS. 5A and 5B, above. In some embodiments, processing unit 110 may analyze the images to determine characteristics associated with an action performed by leading vehicle 200b. For example, where leading vehicle 200b makes a turn, processing unit 110 may determine position information for leading vehicle 200b right before it begins the turn, position information for leading vehicle 200b as it completes the turn, a turning radius, and a speed profile describing changes in speed (e.g., over time) for leading vehicle 200b as it makes the turn.

[0238] At step 1870, processing unit 110 may execute action response module 1706 to cause primary vehicle 200a to mimic the one or more actions of leading vehicle 200b determined at step 1850. For example, processing unit 110 may transmit electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to trigger a turn, lane shift, change in speed, and/or change in direction of primary vehicle 200a. In addition, processing unit 110 may construct a set of path points to guide changes in position and speed for primary vehicle 200a. The path points may be based on lane marks (e.g., midpoint 1610, left edge 1620, right edge 1630, etc.) represented by a polynomial (e.g., of the third-degree) and target coordinates associated with a target vehicle, such as leading vehicle 200b. Processing unit 110 may adjust the coefficients of the polynomial so that the path points pass through the target coordinates. Processing unit 110 may also offset the path points (e.g., relative to left edge 1620 and right edge 1630) based on a predetermined offset amount. As processing unit 110 adjusts the polynomial coefficients and offsets the path points, processing unit 110 may also calculate local curvature information for each of the resulting segments formed between the path points.

[0239] In some embodiments, processing unit 110 (e.g., via execution of action response module 1706) may cause primary vehicle 200a to mimic leading vehicle 200b within certain position, speed, and/or acceleration constraints associated with the two vehicles. For exam-

IPR2025-00943
Tesla EX1002 Page 1708

ple, the constraints may require a predetermined minimum distance (e.g., a safety distance) between primary vehicle 200a and leading vehicle 200b or a predetermined maximum difference in speed and/or acceleration between the vehicles. Thus, if mimicking an action performed by leading vehicle 200b would violate one or more constraints (e.g., mimicking an acceleration would result in primary vehicle 200a being too close to leading vehicle 200b), processing unit 110 may cause primary vehicle 200a to decline to mimic the one or more actions of leading vehicle 200b. In addition, even while mimicking an action performed by leading vehicle 200b (e.g., a lane shift), processing unit 110 may cause primary vehicle 200a to follow leading vehicle 200b's speed profile and maintain a predetermined distance or a safety distance from leading vehicle 200b. Processing unit 110 may accomplish this by causing primary vehicle 200a to: (i) continuously close the distance gap from leading vehicle 200b when the gap is greater than the predetermined minimum distance, (ii) continuously close the gap between the speeds of primary vehicle 200a and leading vehicle 200b, and (iii) match the acceleration of leading vehicle 200b to the extent it is consistent with predetermined constraints. To control the extent to which of these actions should be performed, processing unit 110 may associate each action with weights in real-time based on the current and predicted distance to leading vehicle 200b. Thus, for example, as the distance to leading vehicle 200b decreases more weight may be put on decelerating rather than matching the acceleration of leading vehicle 200b.

[0240] As another example, in the case where primary vehicle 200a is approaching leading vehicle 200b at a speed greater than the speed of leading vehicle 200b, processing unit 110 may cause primary vehicle 200a to gradually slow down such that, when the distance between the two vehicles equals a predetermined safety distance such as a minimum following distance, the speed of the two vehicles is the same. In another example, where a leading vehicle 200b abruptly appears in front of primary vehicle 200a, causing the distance between the two vehicles to be less than the minimum following distance, processing unit 110 may cause primary vehicle 200a to gradually slow down to avoid approaching leading vehicle 200b any closer while at the same increasing the distance between the vehicles until it reaches the minimum following distance.

[0241] FIG. 19 is a flow chart showing an exemplary process 1900 for causing primary vehicle 200a to decline to mimic a turn of leading vehicle 200b, consistent with disclosed embodiments. At step 1910, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 1810 of FIG. 18, above. At step 1920, processing unit 110 may execute action detection module 1704 to locate a leading vehicle 200b within the images received at step 1910 by analyzing the images, as described in connection with step 1850 of FIG. 18, above. At step 1925, processing unit 110 may execute action detection module 1704 to determine whether leading vehicle 200b is turning at an intersection, using the techniques described in connection with step 1860 of FIG. 18, above. If leading vehicle 200b is determined to be turning at the intersection (step 1925, yes), processing unit 110 may cause primary vehicle 200a to decline to mimic the turn of leading vehicle 200b at step 1930. Otherwise (step 1925, no), process 1900 concludes.

[0242] FIG. 20 is a flow chart showing an exemplary process 2000 for causing primary vehicle 200a to mimic or decline a turn of leading vehicle 200b, consistent with disclosed embodiments. At step 2010, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 1810 of FIG. 18, above. At step 2020, processing unit 110 may execute action detection module 1704 to locate a leading vehicle 200b within the images received at step 2010 by analyzing the images, as described in connection with step 1850 of FIG. 18, above. At step 2025, processing unit 110 may execute action detection module 1704 to determine whether leading vehicle 200b is turning at an intersection, using the techniques described in connection with step 1860 of FIG. 18, above. If leading vehicle 200b is determined to be turning at the intersection (step 2025, yes), at step 2035, processing unit 110 may execute action detection module 1704 to determine whether leading vehicle 200b is changing lanes within the intersection, using the techniques described in connection with step 1860 of FIG. 18, above. Otherwise (step 2025, no), process 2000 concludes. If leading vehicle 200b is determined to be changing lanes within the intersection (step 2035, yes), processing unit 110 may cause primary vehicle 200a to decline to mimic the turn of leading vehicle 200b at step 2050. Otherwise (step 2035, no), processing unit 110 may cause primary vehicle 200a to mimic the turn of leading vehicle 200b at step 2040.

[0243] FIG. 21 is a flow chart showing another exemplary process 2100 for causing primary vehicle 200a to mimic one or more actions of a first leading vehicle, consistent with disclosed embodiments. At step 2110, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 1810 of FIG. 18, above. At steps 2120 and 2130, processing unit 110 may execute action detection module 1704 to locate a first and second leading vehicle within the images received at step 2110 by analyzing the images, as described in connection with step 1850 of FIG. 18, above. At step 2135, processing unit 110 may execute action detection module 1704 to determine whether the first leading vehicle path has a lower turning radius than the second leading vehicle path, using the techniques described in connection with step 1860 of FIG. 18, above. If the first leading vehicle path is determined to have a lower turning radius than the second leading

IPR2025-00943
Tesla EX1002 Page 1709

vehicle path (step 2135, yes), processing unit 110 may cause primary vehicle 200a to mimic the turn of the first leading vehicle at step 2140. Otherwise (step 2135, no), process 2100 concludes.

[0244] FIG. 22 is a flow chart showing an exemplary process 2200 for causing primary vehicle 200a to mimic one or more actions of leading vehicle 200b based on a navigation history, consistent with disclosed embodiments. At step 2210, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 1810 of FIG. 18, above. At step 2220, processing unit 110 may execute action detection module 1704 to locate a leading vehicle 200b within the images received at step 2210 by analyzing the images, as described in connection with step 1850 of FIG. 18, above.

[0245] At step 2230, processing unit 110 may determine position information for leading vehicle 200b by analyzing the images received at step 2210. Position information for leading vehicle 200b may indicate the position of leading vehicle 200b relative to primary vehicle 200a. Processing unit 110 may analyze the images to derive position information for leading vehicle 200b and one or more actions taken by leading vehicle 200b by using the techniques described in connection with FIGS. 5A-5D, above.

[0246] At step 2240, processing unit 110 may compare the position information for leading vehicle 200b determined at step 2230 to predetermined map data (e.g., data stored in map database 160). For example, by comparing the position information to predetermined map data, processing unit 110 may derive position coordinates for leading vehicle 200b. In addition, processing unit 110 may use the predetermined map data as an indicator of a confidence level associated with the position information determined at step 2230.

[0247] At step 2250, processing unit 110 may create a navigation history associated with leading vehicle 200b by, for example, tracking the actions taken by leading vehicle 200b. Processing unit 110 may associate the actions taken by leading vehicle 200b (e.g., determined from analysis of the images at step 2230) with predetermined map data from step 2240. Thus, the navigation history may indicate that leading vehicle 200b performed certain actions (e.g., a turn, a lane shift, etc.) at certain position coordinates.

[0248] At step 2260, processing unit 110 may execute action response module 1706 to determine whether to cause primary vehicle 200a to mimic one or more actions of leading vehicle 200b based on the navigation history created at step 2250 and using the techniques described in connection with step 1870 of FIG. 18, above. For example, processing unit 110 (e.g., via action response module 1706) may cause primary vehicle 200a to mimic a turn of leading vehicle 200b, even where no turn is detected from analysis of the images (e.g., at step 2230), when the navigation history indicates that leading vehicle 200b has successfully navigated at least one prior turn. More generally, processing unit 110 may cause primary vehicle 200a to mimic an action taken by leading vehicle 200b in cases where the navigation history indicates the leading vehicle 200b had successfully performed such an action.

Navigating a Vehicle to Pass Another Vehicle

[0249] System 100 may provide driver assist functionality that monitors the vicinity around vehicle 200 and aborts lane changes in the event that the lane change is deemed to be at risk of causing a collision, such as with another vehicle (e.g., a target vehicle). For example, system 100 may enable vehicle 200 to complete a pass of a target vehicle if the target vehicle is determined to be in a different lane than vehicle 200. If, before completion of a pass, system 100 determines that the target vehicle is entering the lane in which vehicle 200 is traveling, system 100 may cause vehicle 200 to abort the pass. For example, aborting the pass may include causing vehicle 200 to brake or causing vehicle 200 to stop accelerating.

[0250] As illustrated in FIG. 23, vehicle 200 may travel on roadway 2300. Vehicle 200 may be equipped with system 100 and may implement any of the disclosed embodiments for identifying lane constraints and operating vehicle 200 within the lane constraints. System 100 may also facilitate lane changes made by vehicle 200. For example, system 100 may determine whether circumstances allow vehicle 200 to safely change lanes and, if the circumstances change during the lane change, cause vehicle 200 to abort lane change.

[0251] As shown in FIG. 23 and as previously discussed with respect to FIG. 8, vehicle 200 may include first vehicle side 802, which may be first distance 805 from a first lane constraint 2302. Similarly, vehicle 200 may include second vehicle side 812 opposite from first vehicle side 802, and second vehicle side 812 may be second distance 815 from a second lane constraint 2304. In this manner, first lane constraint 810 and second lane constraint 2304 may define a lane 2330 within which vehicle 200 may travel.

[0252] System 100 may determine first lane constraint 2302 and second lane constraint 2304 based on a plurality of images acquired by one or more of image capture devices 122-126. According to some embodiments, first lane constraint 2302 and/or second lane constraint 2304 may be identified by visible lane boundaries, including lines marked on a road surface. Additionally or alternatively, first lane constraint 2302 and/or second lane constraint 2304 may include an edge of a road surface. According to some embodiments, system 100 may determine first lane constraint 2302 and/or second lane constraint 2304 by identifying a midpoint of a road surface width 2320 (or any other suitable road feature that vehicle 200 may use as a navigation reference). System 100 may identify lane constraints in this manner when, for example, lines designating road lanes are not painted or

IPR2025-00943
Tesla EX1002 Page 1710

otherwise labeled.

[0253] When passing a target vehicle 2310, the lane change of vehicle 200 may be complicated by a change in the trajectory of target vehicle 2330. Thus, system 100 may identify target vehicle 2310 based on an analysis of a plurality of images acquired by one or more of image capture devices 122-126. System 100 may also identify second lane 2340 in which target vehicle 2310 is traveling. Further, system 100 may determine whether second lane 2340 is different than lane 2330 in which vehicle 200 is traveling. If second lane 2340 is different than lane 2330, system 100 may be configured to enable vehicle 200 to pass target vehicle 2310.

[0254] System 100 may also monitor target vehicle 2310 based on the plurality of images. For example, system 100 may monitor a position of target vehicle 2310, such as monitoring its position relative to vehicle 200, relative to first lane constraint 2302 and/or second lane constraint 2304, an absolute position, or relative to another reference point. Additionally or alternatively, monitoring a position of target vehicle 2310 may include estimating a speed of target vehicle 2310.

[0255] If system 100 determines that target vehicle 2310 is entering lane 2330 in which vehicle 200 is traveling, system 100 may cause vehicle 200 to abort the pass of target vehicle 2310. For example, system 100 may cause vehicle 200 to stop accelerating or brake. Additionally or alternatively, system 100 may issue an audible announcement, such as to notify driver of vehicle 200 that the lane change is being aborted.

[0256] FIG. 24A is an exemplary block diagram of memory 140, which may store instructions for performing one or more operations consistent with disclosed embodiments. As illustrated in FIG. 24A, memory 140 may store one or more modules for identification of the target vehicle and responses described herein. For example, memory 140 may store a lane constraint module 2400, a target vehicle acquisition module 2410, and an action module 2415.

[0257] In one embodiment, lane constraint module 2400 may store instructions which, when executed by processing unit 110, may detect and define first lane constraint 2302 and second lane constraint 2304. For example, processing unit 110 may execute lane offset module 2400 to process the plurality of images received from at least one image capture device 122-124 and detect first lane constraint 2302 and second lane constraint 2304. As discussed above, this may include identifying painted lane lines or markers and/or measuring a midpoint of a road surface. Processing unit 110 may perform the analysis based on the techniques described in connection with FIGS. 5A-5D, above.

[0258] In some embodiments, target acquisition module 2410 may store instructions which, when executed by processing unit 110, may detect the presence of target vehicle 2310 and monitor the position of target vehicle 2310. For example, target acquisition module 2410 may process the plurality of images to detect and monitor target vehicle 2310. Additionally or alternatively, when executing target acquisition module 2410, processing unit 110 may receive information from another module or other system indicative of the presence and/or location of target vehicle 2310. Processing unit 110 may perform the analysis based on the techniques described in connection with FIGS. 5A-5D, above.

[0259] Processing unit 110 may be configured to determine the presence of one or more objects such as target vehicle 2310 in the vicinity of vehicle 200. For example target vehicle 2310 may comprise another vehicle, such as a car, truck or motorcycle traveling near vehicle 200. Processing unit 110 may determine, for each target vehicle 2310, an offset profile. The offset profile may include a determination of whether in its current and predicted position, target vehicle 2310 will be within a predefined range of vehicle 200. If processing unit 110 determines that target vehicle 2310 is or will be within a predefined range of vehicle 200, processing unit may determine whether there is enough space within first lane constraint 810 and second lane constraint 820, for vehicle 200 to bypass target vehicle 2310. If there is not enough space, processing unit 110 may execute a lane change. If there is enough space, processing unit 110 may determine whether there is enough time to bypass target vehicle 2310 before the distance between vehicle 200 and target vehicle 2310 closes. If there is enough time, processing unit 110 may initiate and or activate the offset maneuver to bypass target vehicle 2310. The offset maneuver may be executed by processing unit 110 so that the movement of vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 to 0.75 m/sec. The offset maneuver may be executed so that the maximum amplitude of the offset maneuver will occur when the gap between vehicle 200 and object 825 closes. If there is not enough time, processing unit 110 may initiate braking and then initiate an offset maneuver.

[0260] Further, processing unit 110 may execute action module 2415 to abort the plan of vehicle 200 to pass target vehicle 2310. If target vehicle 2310 is not in lane 2330 of vehicle 200, instructions included in action module 1230 may enable vehicle 200 to pass target vehicle 2330. However, if execution of target vehicle acquisition module 2410 indicates that target vehicle 2310 is changing lanes into lane 2330 of vehicle 200, processing unit 110 may execute instructions included in action module 2415 to abort the lane change of vehicle 200. For example, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to abort the lane change. Additionally or alternatively, processing unit may cause an audible notification to sound.

[0261] FIG. 24B illustrates a flowchart of an exemplary process 2420 for navigating a vehicle among encroaching vehicles, consistent with disclosed embodiments. Process 2420 may identify lane constraints that define a lane of travel for a vehicle (e.g., vehicle 200) to travel,

IPR2025-00943
Tesla EX1002 Page 1711

identify and monitor a target vehicle (e.g., vehicle 2310), and enable the vehicle to pass the target vehicle or abort the pass if it is determined that the target vehicle is entering a lane in which the vehicle is traveling.

[0262] At step 2430, at least one of image capture device 122, 124, and/or 126 may acquire a plurality of images of an area in the vicinity of vehicle 200. For example, a camera included in image acquisition unit 120 (such as image capture device 122) may capture a plurality of images and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive more than one plurality of images via a plurality of data interfaces. For example, processing unit 110 may receive a plurality of images from each of image capture devices 122, 124, 126, each of which may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0263] Next, at step 2440, processing unit 110 may determine a first lane constraint on first vehicle side 802 and a second lane constraint on second vehicle side 812 using the plurality of images received via data interface 128. For example, as part of implementing step 2440, processing unit 110 may execute instructions of lane constraint module 910. Further, as part of determining the first and second lane constraints, processing unit 110 may use one or more of the processes discussed above in connection with FIGS. 5A-5D.

[0264] At step 2450, processing unit 110 may acquire target vehicle 2310. To acquire target vehicle 2310, processing unit 110 may execute target vehicle acquisition module 2410. For example, processing unit 110 may execute instructions to analyze acquired images to identify second lane 2340 in which target vehicle 2310 is traveling and determine whether target vehicle 2310 is traveling in lane 2330 in which vehicle 200 is traveling. This determination may be based on the techniques described in connection with FIGS. 5A-5D, above.

[0265] At step 2455, processing unit 110 may determine whether second lane 2340 in which target vehicle 2310 is traveling is different than lane 2330 in which vehicle 200 is traveling.

[0266] At step 2460, processing unit 110 may enable vehicle 200 to pass target vehicle 2310 if second lane 2340 is different than lane 2330 in which vehicle 200 is traveling. To enable vehicle 200 to pass target vehicle 2310, processing unit may execute action module 2415.

[0267] Next, at step 2465, processing unit 110 may execute target acquisition module 2410 to monitor target vehicle 2310. This may include estimating a velocity of target vehicle 2310 and determining the location of target vehicle 2310, such as with respect to vehicle 200.

[0268] Step 2470 may include determining whether target vehicle 2310 has entered lane 2330. If processing unit 110 determines that target vehicle 2310 has not entered lane 2330, at step 2480, processing unit 110 may allow vehicle 200 to complete the pass.

[0269] At step 2490, if processing unit 110 determines that target vehicle 2310 has entered lane 2330 (or is entering lane 2330 or is otherwise on a trajectory to bring vehicle 2310 into lane 2330), processing unit 110 may execute action module 2415 to abort the pass. This may include, for example, issuing an audible notification to driver of vehicle 200 via speakers 360. Additionally or alternatively, step 2490 may include executing instructions included in action module 2415 to cause processing unit 110 to transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200 to abort lane change.

Navigating a Vehicle to Avoid Encroaching Traffic

[0270] System 100 may provide driver assist functionality that monitors the vicinity around vehicle 200 and responds to the presence of another vehicle encroaching upon vehicle 200. Monitoring the vicinity around vehicle 200 may include monitoring a collision threshold. For example a collision threshold may include a time to collision or a minimum predetermined distance between vehicle 200 and other traffic. System 100 may further cause vehicle 200 to take evasive action to avoid encroaching traffic. For example, evasive action may include braking, changing lanes, or otherwise altering the course of vehicle 200.

[0271] As illustrated in FIG. 25, vehicle 200 may travel on roadway 2500. Vehicle 200 may include system 100 to detect and respond to encroaching traffic. For example, vehicle 200 may determine that another vehicle is entering into the lane in which vehicle 200 is traveling or otherwise determine that another vehicle is crossing a collision threshold. Vehicle 200 may react to an encroaching vehicle, such as by altering course, braking, or accelerating.

[0272] As discussed with respect to FIG. 8 and as shown in FIG. 25, vehicle 200 may include first vehicle side 802, which may be first distance 805 from a first lane constraint 2502. Similarly, vehicle 200 may include second vehicle side 812 opposite from first vehicle side 802, and second vehicle side 812 may be second distance 815 from a second lane constraint 2504. In this manner, first lane constraint 2502 and second lane constraint 2504 may define a lane 2505 within which vehicle 200 may travel.

[0273] System 100 may determine first lane constraint 2502 and second lane constraint 2504 based on a plurality of images acquired by one or more of image capture devices 122-126. According to some embodiments, first lane constraint 2502 and/or second lane constraint 2504 may be identified by visible lane boundaries, such a line marked on a road surface. Additionally or alternatively, first lane constraint 2502 and/or second lane constraint 2504 may include an edge of a road surface. According to some embodiments, system 100 may determine first lane constraint 2502 and/or second lane constraint 2502 by identifying a midpoint 2520 of a road surface width (or

IPR2025-00943
Tesla EX1002 Page 1712

any other suitable road feature vehicle 200 may use as a navigation reference). System 100 may identify lane constraints in this manner when, for example, lines designating road lanes are not painted or otherwise labeled.

[0274] Detection of first lane constraint 2502 and/or second lane constraint 2504 may include processing unit 110 determining their 3D models in a camera coordinate system. For example, the 3D models of first lane constraint 2502 and/or second lane constraint 2504 may be described by a third-degree polynomial. In addition to 3D modeling of first lane constraint 2502 and/or second lane constraint 2504, processing unit 110 may estimate motion parameters, such as the speed, yaw and pitch rates, and acceleration of vehicle 200. Optionally, processing unit may detect static and moving vehicles and their position, heading, speed, and acceleration, all relative to vehicle 200. Processing unit 110 may determine a road elevation model to transform the information acquired from the plurality of images into 3D space.

[0275] The driver of vehicle 200 cannot assume that other traffic traveling in adjacent or near-adjacent lanes will only change lanes if it is safe for that traffic to shift into lane 2505 in which vehicle 200 is driving. Thus, it may be desirable to consider and monitor traffic traveling in nearby lanes to detect lateral encroachment into lane 2505 of vehicle 200, so that vehicle 200 may react to lateral encroachment to avoid a collision. System 100 may identify an encroaching vehicle 2510 based on the plurality of images and determine that encroaching vehicle 2510 is approaching vehicle 200. System 100 may then cause an action to account for the encroaching vehicle. For example, system 100 may cause vehicle 200 to maintain a current velocity and travel within first lane constraint 2502 and second lane constraint 2504 such that a first distance 2530, which is on first vehicle side 802 on which encroaching vehicle 2510 is approaching, is greater than a second distance 2540.

[0276] Processing unit 110 may be configured to determine the presence of one or more encroaching vehicles 2510 in the vicinity of vehicle 200. Processing unit 110 may determine, for each encroaching vehicle 2510 an offset profile. The offset profile may include a determination of whether in its current and predicted position, encroaching vehicle 2510 will be within a predefined range of vehicle 200. If processing unit 110 determines that encroaching vehicle 2510 is or will be within a predefined range of vehicle 200, processing unit may determine whether there is enough space within first lane constraint 2502 and second lane constraint 2504, for vehicle 200 to bypass encroaching vehicle 2510. If there is not enough space, processing unit 110 may execute a lane change. If there is enough space, processing unit 110 may determine whether there is enough time to bypass encroaching vehicle 2510 before the distance between vehicle 200 and encroaching vehicle 2510 closes.

[0277] For example, if there is enough time, processing unit 110 may initiate and or activate the offset maneuver to bypass encroaching vehicle 2510. The offset maneuver may be executed by processing unit 110 so that the movement of vehicle 200 is smooth and comfortable for the driver. For example, the offset slew rate may be approximately 0.15 to 0.75 m/sec. The offset maneuver may be executed so that the maximum amplitude of the offset maneuver will occur when the gap between vehicle 200 and encroaching vehicle 2510 closes. If there is not enough time, processing unit 110 may initiate braking (e.g., via transmitting electronic signals to braking system 230) and then initiate offset maneuver (e.g., by transmitting electronic signals to throttling system 220 and/or steering system 240).

[0278] System 100 may determine, based on a plurality of images acquired by one or more of image capture devices 122-126, that encroaching vehicle 2510 has crossed at least one collision threshold. For example, the at least one collision threshold may include a minimum predetermined distance 2550 between vehicle 200 and encroaching vehicle 2510. For example, encroaching vehicle 2510, as illustrated in FIG. 25, is shown exceeding minimum predetermined distance 2550. Additionally or alternatively, the at least one collision threshold may include a time to collision of vehicle 200 and encroaching vehicle 2510. The time-to-collision threshold may be determined based on the distance between vehicle 200 and encroaching vehicle 2510, the velocity of vehicle 200 and/or encroaching vehicle 2510, the lateral velocity of vehicle 200 and/or encroaching vehicle 2510, and/or the acceleration of vehicle 200 and/or encroaching vehicle 2510. A collision threshold based on a time to collision may be set to allow enough time for system 100 to effectively conduct evasive action to avoid collision.

[0279] If system 100 determines that at least one collision threshold has been crossed, system 100 may cause vehicle 200 to take evasive action. For example, this may include causing other subsystems of vehicle 200 to operate. Thus, system 100 may operate to cause vehicle 200 to increase speed by communicating with throttling system 220, braking system 230, and/or steering system 240 to change the speed and/or direction of vehicle 200.

[0280] The evasive action may include causing vehicle 200 to change lanes. This may be a desirable response if, for example, vehicle 200 is traveling on a multilane highway and system 100 determines that changing lanes to move away from encroaching vehicle 2510 will avoid a collision. Evasive action may include vehicle 200 altering its course. Additionally or alternatively, evasive action may include altering course of vehicle 200. For example, to take evasive action, processing unit 110 may transmit electronic signals to throttling system 220, braking system 230, and/or steering system 240.

[0281] System 100 may also be configured to cause vehicle 200 to take an action after the evasive action has been completed. For example, system 100 may cause vehicle 200 to return to midpoint 2520 of lane 2505 in which vehicle 200 is traveling. Additionally or alternatively, system 100 may cause vehicle 200 to resume a preset

IPR2025-00943
Tesla EX1002 Page 1713

speed after braking in response to encroaching vehicle 2510.

[0282] FIG. 26 is an exemplary block diagram of memory 140, which may store instructions for detecting and responding to traffic laterally encroaching on a vehicle consistent with disclosed embodiments. As illustrated in FIG. 26, memory 140 may store a lane constraint module 2610, an encroaching vehicle acquisition module 2620, and an action module 2630.

[0283] In one embodiment, lane constraint module 2610 may store instructions which, when executed by processing unit 110, may detect and define first lane constraint 2502 and second lane constraint 2504. For example, processing unit 110 may execute lane offset constraint module 2610 to process the plurality of images received from at least one image capture device 122-126 and detect first lane constraint 2502 and second lane constraint 2504. As discussed above, this may include identifying painted lane lines and/or measuring a midpoint of a road surface. Processing unit 110 may perform the analysis based on the techniques described in connection with FIGS. 5A-5D, above.

[0284] In one embodiment, encroaching vehicle acquisition module 2620 may store instructions which, when executed by processing unit 110, may detect the presence of encroaching vehicle 2510 and monitor encroaching vehicle 2510. For example, encroaching vehicle acquisition module 2620 may include instructions for determining a relative location of encroaching vehicle 2510 to vehicle 200. In some embodiments, this may include monitoring the velocity of encroaching vehicle 200. According to some embodiments, encroaching vehicle acquisition module 2620 may determine whether encroaching vehicle 2510 has exceeded a collision threshold. A collision threshold may be, for example, a minimum predetermined distance between vehicle 200 and encroaching vehicle 2510, If encroaching vehicle 2510 comes closer to vehicle 200 than the minimum predetermined distance, processing unit 110 may determine that encroaching vehicle 2510 is indeed encroaching and execute instructions to avoid collision with encroaching vehicle 2510.

[0285] Additionally or alternatively, a collision threshold may include a time to collision. The time-to-collision threshold may be determined based on the distance between vehicle 200 and encroaching vehicle 2510, the velocity of vehicle 200 and/or encroaching vehicle 2510, the lateral velocity of vehicle 200 and/or encroaching vehicle 2510, and/or the acceleration of vehicle 200 and/or encroaching vehicle 2510. A collision threshold based on a time to collision may be set to allow enough time for processing unit 110 and system 100 to effectively conduct evasive action to avoid collision.

[0286] Further, processing unit 110 may execute action module 2630 to respond to the detection of encroaching vehicle 2510 and/or the crossing of a collision threshold. For example, action module 2630 may include instructions to cause vehicle 200 to maintain a current velocity and to travel within first lane constraint 2502 and second lane constraint 2504 such that first distance 2530 is greater than second distance 2540. The first distance 2530 may be on the same side of vehicle 200 as encroaching vehicle 2510, so that vehicle 200 is farther away from encroaching vehicle 2510 while still within first lane constraint 2502 and second lane constraint 2504. If processing unit 110 determines that a collision threshold has been crossed, processing unit 110 may execute action module 2630 to cause vehicle 200 to take evasive action, such as altering course, braking, and/or changing lanes. Action module 2630 may also include instructions that cause vehicle 200 to resume course after completing evasive action. For example, processing unit 110 may cause vehicle 200 to resume a preset speed after braking or return to the center of lane 2505.

[0287] FIG. 27 illustrates a flowchart of an exemplary process 2720 detecting and responding to traffic laterally encroaching on a vehicle, consistent with disclosed embodiments. Process 2720 may identify lane constraints that define a lane of travel for vehicle travel, determine whether an encroaching vehicle 2510 is approaching, and cause vehicle 200 to maintain current velocity and to travel within first lane constraint 2501 and second lane constraint 2504 such that first distance 2530, on the side of vehicle 200 that encroaching vehicle 2510 is approaching, is greater than second distance 2540.

[0288] At step 2730, at least one image capture device 122,124, and/or 126 may acquire a plurality of images of an area in the vicinity of vehicle 200. For example, a camera included in image acquisition unit 120 (such as image capture device 122) may capture a plurality of images and transmit them over a digital connection (e.g., USB, wireless, Bluetooth, etc.) to processing unit 110. In some embodiments, processing unit 110 may receive more than one plurality of images via a plurality of data interfaces. For example, processing unit 110 may receive a plurality of images from each of image capture devices 122, 124, 126, each of which may have an associated data interface for communicating data to processing unit 110. The disclosed embodiments are not limited to any particular data interface configurations or protocols.

[0289] Next, at step 2740, processing unit 110 may determine a first lane constraint on first vehicle side 802 and a second lane constraint on second vehicle side 812 using the plurality of images received via data interface 128. For example, as part of implementing step 2740, processing unit 110 may execute instructions of lane constraint module 2610. Further, as part of determining the first and second lane constraints, processing unit 110 may use one or more of the processes discussed above in connection with FIGS. 5A-5D.

[0290] At step 2750, processing unit 110 may determine whether encroaching vehicle 2510 is approaching. To make this determination, processing unit 110 may execute encroaching vehicle acquisition module 2620. Encroaching vehicle acquisition module 2620 may include instructions to process images to detect other traffic. This

IPR2025-00943
Tesla EX1002 Page 1714

may include identifying vehicles within a certain range of vehicle 200, such as those vehicles that are adjacent to vehicle 200. These adjacent vehicles may be encroaching vehicles 2510.

[0291] Once encroaching vehicle 2510 is identified, processing unit may determine more information regarding vehicle 2510. For example, at step 2555, processing unit 110 may determine whether target vehicle 2330 has exceeded a collision threshold. A collision threshold may be, for example, a minimum predetermined distance between vehicle 200 and encroaching vehicle 2510. If encroaching vehicle 2510 comes closer to vehicle 200 than the minimum predetermined distance, processing unit 110 may determine that encroaching vehicle 2510 is indeed encroaching and execute instructions to avoid collision with encroaching vehicle 2510.

[0292] Additionally or alternatively, a collision threshold may be a time to collision. The time-to-collision threshold may be determined based on the distance between vehicle 200 and encroaching vehicle 2510, the velocity of vehicle 200 and/or encroaching vehicle 2510, the lateral velocity of vehicle 200 and/or encroaching vehicle 2510, and/or the acceleration of vehicle 200 and/or encroaching vehicle 2510. A collision threshold based on a time to collision may be set to allow enough time for processing unit 110 and system 100 to effectively conduct evasive action to avoid collision.

[0293] At step 2760, processing unit 110 may cause vehicle 200 to maintain current velocity and to travel within first lane constraint 2502 and second lane constraint 2504. To cause vehicle 200 to operate in this manner, processing unit 110 may execute action module 2630. Additionally or alternatively, step 2760 may include executing instructions on action module 1230 to cause processing unit 110 to transmit electronic signals to the accelerator 2610, brakes 2620, and/or steering system of vehicle 200 to cause vehicle 200 to maintain current velocity and to travel within first and second lane constraints 810 and 820.

[0294] If, at step 2755, processing unit 110 determined that a collision threshold has been exceeded, the processing unit 110 may take evasive action. The evasive action may include causing vehicle 200 to change lanes. This may be a desirable response if, for example, vehicle 200 is traveling on a multilane highway and system 100 determines that changing lanes to move away from encroaching vehicle 2510 will avoid a collision. Evasive action may include vehicle 200 altering its course. Additionally or alternatively, evasive action may include altering course of vehicle 200. For example, processing unit 110 may undertake evasive action by transmitting electronic signals to throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. After the evasive action is completed, at step 2780, processing unit 110 may control vehicle 200 to resume its prior navigation (e.g., to resume navigation to a desired location).

Multi-Threshold Reaction Zone for Vehicle Navigation

[0295] System 100 may provide driver assist functionality that causes a response in vehicle 200, such as braking, accelerating, switching lanes, turning, and/or other navigational responses. For example, if system 100 detects the presence of another vehicle ahead of vehicle 200 that is traveling at a rate of speed slower than vehicle 200, system 100 may cause vehicle 200 to reduce its speed by braking. System 100 may apply a braking profile that accounts for the immediacy of the risk (e.g., the risk of collision) to provide the driver with a natural driving sensation.

[0296] System 100 may cause various different responses depending on the situation and a determined risk of collision with another vehicle, for example. System 100 may take no action in response to changes in speed or direction of a target vehicle ahead if the target vehicle is sufficiently far ahead and the determined risk of collision is low. Where the target vehicle may be closer or when the risk of collision rises above a predetermined level, system 100 may cause a response, such as braking, etc. When the risk of collision is even higher (e.g., where system 100 determines that a collision would be imminent without evasive action), system 100 may cause vehicle 100 to take an evasive action, such as maximum or near maximum braking, change of direction, etc.

[0297] FIGS. 28A and 28B are diagrammatic representations of a primary vehicle 200a on a road 2800, consistent with disclosed embodiments. In some instances, road 2800 may be divided into two lanes, lane 2810 and lane 2820. A target object, such as target vehicle 200b, may be traveling ahead of primary vehicle 200a in the same lane (e.g., lane 2820 as shown in FIGS. 28A and 28B) or in a different lane. As shown in the example depicted in FIG. 28A, primary vehicle 200a and target vehicle 200b may be separated by a distance d_1. At a given time after the scene depicted in FIG. 28A, primary vehicle 200a and target vehicle 200b may be separated by a distance d_2 as shown in FIG. 28B, where d_2 is less than d_1. For example, primary vehicle 200a may be traveling at a rate of speed greater than target vehicle 200b, causing the distance between the vehicles to decrease from d_1 to d_2.

[0298] As another example, primary vehicle 200a and target vehicle 200b may be traveling at the same rate of speed until target vehicle 200b reduces its speed by braking, or primary vehicle 200a increases its speed by accelerating, causing the distance between the vehicles to decrease from d_1 to d_2. As another example, target vehicle 200b may encroach from a side of primary vehicle 200a. In some circumstances, if target vehicle 200b is ahead or a sufficient distance away from primary vehicle 200a (e.g., five meters or more, 10 meters or more, etc.) and moves toward primary vehicle 200a, system 100 may not take any action. However, if target vehicle 200b continues to move closer to primary vehicle 200a, then system 100 may cause primary vehicle 200a to brake,

IPR2025-00943
Tesla EX1002 Page 1715

change speed, and/or change lanes. For example, if target vehicle 200b is within a predetermined threshold (e.g., within five meters), primary vehicle 200a may then take action. Different scenarios involving the relative speeds and positions of primary vehicle 200a and target vehicle 200b are possible, and the disclosed embodiments are not limited to any particular scenario.

[0299] In some embodiments, the multi-threshold reaction zone may provide a natural feeling driving experience to the user. For example, as discussed above, if the risk of collision is low, then system 100 may disregard movements by target vehicle 200b. As a result, primary vehicle 200a may not respond to movements or actions of target vehicle 200b when there is not a risk of a collision (e.g., primary vehicle 200a would not need to mimic every motion of target vehicle 200b when it's 100 meters ahead). However, if target vehicle 200b is close to primary vehicle 200a and/or if system 100 determines that the risk of collision is high, then a responsive maneuver may be taken such that the user feels safer by increasing space between primary vehicle 200a and target vehicle 200b. If system 100 determines that a crash may be imminent, then system 100 may take a more dramatic action without causing the user to be surprised, because the action may be necessary to avoid a collision.

[0300] Accordingly, in some embodiments, system 100 may associate the different reaction zones with different degrees of responsive actions. For example, in some embodiments, system 100 may take no action when target vehicle 200b is within a first threshold (e.g., when target vehicle 200b is far away from primary vehicle 200a), system 100 may take a medium action with target vehicle 200b is within a second threshold (e.g., when target vehicle 200b is becoming closer to primary vehicle 200a), and system 100 may take a more dramatic action when target vehicle 200 is within a third threshold (e.g., when target vehicle 200b is sufficiently close to primary vehicle 200a that there is a risk of a collision). In a first reaction zone associated with the first threshold, system 100 may take no action when the estimated distance-to-target associated with target vehicle 200b is greater than twice a predetermined safety distance. In some embodiments, the first reaction zone may apply where the estimated distance-to-target associated with target vehicle 200b falls in the range of 1.5 to 2.5 times the predetermined safety distance.

[0301] The estimated distance-to-target may be the current distance between primary vehicle 200a and target vehicle 200b plus the expected change of this distance over a prediction time (e.g., in the range of 0.1 to 1.0 seconds, depending on factors associated with system 100 such as velocity tracking control loops, actuators, and vehicle dynamics). The expected change of distance may be calculated as: $(v_{target}-v_{primary})*t_{prediction} + ((a_{target}-a_{primary})*t_{prediction}^{1/2})$, where v_{target} and $v_{primary}$ correspond to the speed of target vehicle 200b and primary vehicle 200a, respectively, a_{target} and $a_{primary}$ correspond to the acceleration of target vehicle 200b and primary vehicle 200a, respectively, and $t_{prediction}$ corresponds to the prediction time. Choosing a prediction time approximating the overall reaction time of system 100 may account for the fact that desired navigational responses for primary vehicle 200a may not be implemented immediately, and thereby may provide the driver and/or passengers of primary vehicle 200a with a smooth, natural driving experience.

[0302] The predetermined safety distance may be defined as: max $\{d_{static}, d_{dynamic}\}$. d_{static}, which may fall in the range of 2 to 10 meters, may represent a desired safety distance between primary vehicle 200a and another vehicle (such as target vehicle 200b) while the vehicles are stopping and/or moving at a low rate of speed. $d_{dynamic}$ may represent a desired safety distance when the vehicles are moving at a speed greater than a low rate of speed, and may be calculated as: $t_{safety}*\min(v_{target}, v_{primary})$, where v_{target} and $v_{primary}$ correspond to the speed of target vehicle 200b and primary vehicle 200a, respectively. t_{safety} may fall in the range of 0.5 to 2.5 seconds and may be adjusted by the driver of primary vehicle 200a via user interface 170 according to the driver's preferences. Thus, the driver may be able to control the distance that system 100 maintains between primary vehicle 200a and other vehicles.

[0303] In a second reaction zone associated with the second threshold, system 100 may place target vehicle 200b in a safety zone and apply different weights to different maneuvers performed by target vehicle 200b. For example, as target vehicle 200b may approach a lower boundary of the safety zone, e.g., the estimated distance-to-target associated with target vehicle 200b may approach the predetermined safety distance (discussed above in connection with the first reaction zone), system 100 may assign more weight to deceleration maneuvers (e.g., mimicked more) and less weight to acceleration maneuvers (e.g., mimicked less). In some embodiments, the lower boundary may fall in the range of 0.5 to 1.5 times the predetermined safety distance, depending on driver preferences. Conversely, as target vehicle 200b approaches an upper boundary of the safety zone, e.g., the estimated distance-to-target associated with target vehicle 200b may approach twice the predetermined safety distance, system 100 may assign less weight to deceleration maneuvers (e.g., mimicked less) and more weight to acceleration maneuvers (e.g., mimicked more). In some embodiments, the lower boundary may fall in the range of 1.5 to 2.5 times the predetermined safety distance, depending on driver preferences.

[0304] In a third reaction zone associated with the third threshold, system 100 may place target vehicle 200b in a danger zone when the estimated distance-to-target associated with target vehicle 200b is less than the predetermined safety distance. In this scenario, system 100 may cause primary vehicle 200a to mimic (e.g., immediately) any deceleration maneuvers (while disregarding any acceleration maneuvers) performed by target vehicle

IPR2025-00943
Tesla EX1002 Page 1716

200b. In some embodiments, the first reaction zone may apply where the estimated distance-to-target associated with target vehicle 200b falls in the range of 0.5 to 1.5 times the predetermined safety distance, depending on driver preferences.

[0305] In some embodiments, system 100 may cause a response in primary vehicle 200a, such as braking. As indicated in FIG. 28B, for example, the distance between the primary vehicle 200a and target vehicle 200b may decrease from d_1 to d_2. Depending on the distance d_2 and, by extension, the time before collision, system 100 may cause primary vehicle 200a to reduce its speed so as to avoid a collision with target vehicle 200b. In addition, system 100 may apply one or more braking profiles when causing primary vehicle 200a to reduce its speed, also based on, for example, the time before collision. Thus, if the time before collision (and/or the distance d_2) exceeds a first predetermined threshold, system 100 may cause primary vehicle 200a to brake in a gradual manner, change lanes, and/or change speed, whereas if the time before collision (and/or the distance d_2) is below a second predetermined threshold, system 100 may cause primary vehicle 200a to brake, change lanes, and/or change speed in a more rapid manner.

[0306] FIG. 29 is an exemplary block diagram of a memory configured to store instructions for performing one or more operations, consistent with the disclosed embodiments. As shown in FIG. 29, memory 140 may store a target monitoring module 2902, an intercept time module 2904, and an action response module 2906. The disclosed embodiments are not limited to any particular configuration of memory 140. Further, processing unit 110 may execute the instructions stored in any of modules 2902-2906 included in memory 140.

[0307] In one embodiment, target monitoring module 2902 may store software instructions (such as computer vision software) which, when executed by processing unit 110, identifies and monitors a target object (such as target vehicle 200b) appearing in a plurality of images representing an environment in the vicinity of primary vehicle 200. For purposes of this disclosure, the target object may also be referred to as target vehicle 200b, although the target object may be a pedestrian, road hazard (e.g., large debris), and the like. The plurality of images may be acquired by one or more of image capture devices 122, 124, and 126. Processing unit 110 may execute target monitoring module 2902 to identify and monitor target vehicle 200b by, for example, analyzing the plurality of images using the techniques described in connection with FIGS. 5A-5D, above. Thus, target monitoring module 2902 may track different variables, such as the motion (e.g., direction of travel and trajectory), speed, and acceleration associated with target vehicle 200b, as well as the distance between primary vehicle 200a and the target vehicle 200b. Such information may be tracked continuously or at predetermined intervals.

[0308] In one embodiment, intercept time module 2904 may store software instructions which, when executed by processing unit 110, determines an indicator of an intercept time between primary vehicle 200a and target vehicle 200b. The intercept time may refer to an amount of time before primary vehicle 200a makes contact (e.g., collides) with target vehicle 200b. Processing unit 110 may execute intercept time module 2904 to determine an indicator of the intercept time based on an analysis of images acquired by one or more of image capture devices 122, 124, and 126. For example, intercept time module 2904 may make the determination using a series of time-based observations, e.g., of the position, speed, and/or acceleration of target vehicle 200b (relative to primary vehicle 200a), such as Kalman filters or linear quadratic estimation (LQE). The Kalman filters may be based on a measurement of the scale of target vehicle 200b, where the scale measurement is proportional to the intercept time.

[0309] In one embodiment, action response module 2906 may store software instructions (such as computer vision software) which, when executed by processing unit 110, causes one or more responses in primary vehicle 200a. The responses may be based on the identification and monitoring of target vehicle 200b (e.g., performed via execution of target monitoring module 2902) and/or the indicator of an intercept time between primary vehicle 200a and target vehicle 200b (e.g., determined via execution of intercept time module 2904). Processing unit 110 may cause one or more responses in primary vehicle 200a by, for example, transmitting electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to brake, accelerate, or trigger a turn, lane shift, and/or change in direction of primary vehicle 200a. Processing unit 110 may use extended Kalman Filters to estimate data associated with target vehicle 200b and primary vehicle 200a, such as: speed, acceleration, turn rate, pitch angle, pitch rate, position and heading of target vehicle 200b relative to primary vehicle 200a, the size of target vehicle 200b, road geometry, and the like. For example, processing unit 110 may receive measurements from components on primary vehicle 200a, such as an odometer, an inertial measurement unit (IMU), image processor 190, and the like. Thus, processing unit 110 may calculate data associated with target vehicle 200b and/or primary vehicle 200a using state equations derived from kinematic equations of motion, rigidity constraints, and motion assumptions, where inputs to the equations may include measurements from components on primary vehicle 200a. Processing unit 110 may also use the techniques described in connection with FIGS. 4-7, above, to cause one or more responses in primary vehicle 200a.

[0310] FIG. 30 is a flow chart showing an exemplary process 3000 for causing a response in a primary vehicle, consistent with disclosed embodiments. At step 3010, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120. For instance, one or more cam-

IPR2025-00943
Tesla EX1002 Page 1717

eras included in image acquisition unit 120 may capture a plurality of images of an area forward of primary vehicle 200a (or to the sides or rear of a vehicle, for example) and transmit them over a data connection (e.g., digital, wired, USB, wireless, Bluetooth, etc.) to processing unit 110.

[0311] At step 3020, processing unit 110 may execute target monitoring module 2902 to identify a target object (such as target vehicle 200b) within the plurality of images. For example, target monitoring module 2902 may detect the presence of vehicles in the images using the techniques described in connection with FIGS. SA and 5B, above. Multiple vehicles (and/or objects) may be detected in the same lane or in a different lane as primary vehicle 200a, and at varying distances from primary vehicle 200a, based on image data, position data (e.g., GPS location information), map data, the yaw rate of primary vehicle 200a, lane or other road markings, speed data, and/or data from sensors included in primary vehicle 200a. In a scenario where multiple vehicles are detected, target monitoring module 2902 may determine target vehicle 200b to be the vehicle traveling in the same lane as primary vehicle 200a and/or located closest to primary vehicle 200a.

[0312] At step 3030, processing unit 110 may execute target monitoring module 2902 to monitor target vehicle 200b identified at step 3020. For example, target monitoring module 2902 may track information associated with target vehicle 200b, such as the motion (e.g., direction of travel and trajectory), speed, and acceleration, as well as the distance between primary vehicle 200a and the target vehicle 200b. Target monitoring module 2902 may track such information continuously, or at predetermined intervals, by analyzing the plurality of images using the techniques described in connection with FIGS. 5A-5D, above. At step 3040, processing unit 110 may execute intercept time module 2904 to determine an indicator of an intercept time between primary vehicle 200a and target vehicle 200b. Intercept time module 2904 may make the determination by tracking the position, speed, and/or acceleration of target vehicle 200b (relative to primary vehicle 200a) over one or more periods of time. Intercept time module 2904 may use the techniques described in connection with FIGS. 5A and 5B, above to track such information.

[0313] At step 3050, processing unit 110 may execute action response module 2906 to cause one or more responses in primary vehicle 200a. For example, processing unit 110 may transmit electronic signals (e.g., via a CAN bus) to throttling system 220, braking system 230, and/or steering system 240 of primary vehicle 200a to brake, accelerate, or trigger a turn, lane shift, and/or change in direction of primary vehicle 200a. Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240. For example, processing unit 110 may transmit electronic signals that cause system 100 to physically depress the brake by a predetermined amount or ease partially off the accelerator of vehicle 200. Further, processing unit 110 may transmit electronic signals that cause system 100 to steer vehicle 200 in a particular direction. Such responses may be based on the monitoring of target vehicle 200b (performed at step 3030) and/or the indicator of an intercept time between primary vehicle 200a and target vehicle 200b (determined at step 3040). In some embodiments, the intercept time may be compared with one or more predetermined thresholds, e.g., a first and second intercept threshold. For example, the response may include an emergency avoidance action if the indicator of intercept time indicates that a target object is within a particular intercept threshold (e.g., within the second intercept threshold). The emergency avoidance action may include changing lanes and/or emergency braking.

[0314] The thresholds may be expressed as units of time. Thus, for example, if the indicator of intercept time indicates that target vehicle 200b is outside of the first threshold (e.g., the time before collision between primary vehicle 200a and target vehicle 200b exceeds the first threshold), then action response module 2906 may cause a response in primary vehicle 200a based on an averaged motion profile associated with target vehicle 200b. The averaged motion profile associated with target vehicle 200b may represent an averaged position and/or averaged speed of target vehicle 200b for a particular threshold. In some embodiments, if target vehicle 200b is a certain distance and/or traveling a certain speed, system 100 may elect to take no action. For example, if target vehicle 200b is ahead or a sufficient distance away from primary vehicle 200a (e.g., five meters, 10 meters, etc.), system 100 may not take any action unless target vehicle 200b becomes closer to primary vehicle 200.

[0315] If the indicator of intercept time indicates that target vehicle 200b is between the first and second thresholds (e.g., the time before collision between primary vehicle 200a and target vehicle 200b falls in between the first and second thresholds), then action response module 2906 may cause a response in primary vehicle 200a based on the actual motion associated with target vehicle 200b. The actual motion may represent a position and/or speed of target vehicle 200b (e.g., not averaged or delayed) If the indicator of intercept time indicates that target vehicle 200b is within the second threshold (e.g., the time before collision between primary vehicle 200a and target vehicle 200b is below the second threshold), then action response module 2906 may cause an emergency action in primary vehicle 200a (such as a lane change or emergency braking).

[0316] FIG. 31 is a flow chart showing an exemplary process 3100 for declining to cause a response in a primary vehicle, consistent with disclosed embodiments. At step 3110, processing unit 110 may receive a plurality of images via data interface 128 between processing unit 110 and image acquisition unit 120, as described in connection with step 3010 of FIG. 30, above.

[0317] At step 3120, processing unit 110 may execute

IPR2025-00943
Tesla EX1002 Page 1718

target monitoring module 2902 to locate a target object (e.g., target vehicle 200b) within the images received at step 3010 by analyzing the images, as described in connection with step 3020 of FIG. 30, above. At step 3125, processing unit 110 may execute target monitoring module 2902 to determine whether the target object (e.g., target vehicle 200b) is traveling in a lane different from primary vehicle 200a, using the techniques described in connection with step 3030 of FIG. 30, above. If target vehicle 200b is determined to be traveling in a lane different from primary vehicle 200a (step 3125, yes), processing unit 110 may decline to cause a response in primary vehicle 200a at step 3130. Alternatively, if system 100 determines that target vehicle 200b does not constitute an imminent risk of collision (e.g., target vehicle 200b a nonintersecting trajectory with primary vehicle 200a or a potential time to collision beyond a threshold such that the target object does not pose an imminent risk at the present time), processing unit 110 may decline to cause a response in primary vehicle 200a at step 3130. Otherwise (step 3125, no), process 3100 may conclude and no response occurs.

Stopping an Autonomously Driven Vehicle Using a Multi-Segment Braking Profile

[0318] System 100 may provide driver assist functionality that monitors and responds to environmental conditions to control the braking of vehicle 200. For instance, system 100 may detect, based on a plurality of images, an object that triggers the stopping of vehicle 200. For example, this object may include another vehicle in the vicinity of vehicle 200, a traffic light, and/or a traffic sign. System 100 may use a braking profile to slow and/or stop vehicle 200. According to some embodiments, the braking profile may include multiple segments, where deceleration is executed at different stages. For example, the braking profile may cause system 100 to stop vehicle 200 at a decreasing rate of deceleration.

[0319] By providing a braking profile that includes multiple segments, system 100 may cause vehicle 200 to operate in a manner that feels natural to a driver and/or passengers. For example, when a driver stops a vehicle at a traffic light, the brakes are typically not applied at a constant rate from the moment they are applied until the moment when the vehicle stops. Operating brakes in such a manner result in an abrupt stop at the end. Instead, the driver may apply the brakes gently at first, then with increasing pressure, and then as the vehicle reaches its target stopping location, the driver may decrease the level of braking until it reaches zero (or near zero) at the moment the car stops. Operating the vehicle in this manner may provide a natural stopping progression, i.e., the driver may cause-a progressive deceleration at first so the vehicle does not experience an abrupt initial deceleration and then at the end, so that vehicle does not come to an abrupt stop.

[0320] System 100 may provide autonomous vehicle navigation that may implement braking according to an approach that feels natural to the driver and/or passengers of vehicle 200. For example, in some embodiments, system 100 may cause one or more of image capture devices 122-126 to acquire a plurality of images of an area in a vicinity of vehicle 200. Processing device 110 may receive the receive the plurality of images via data interface 128 and identify, based on analysis of the plurality of images, a trigger (e.g., an object, a traffic light, and/or a traffic sign) for stopping vehicle 200. Based on the identified trigger, system 100 may cause vehicle 200 to stop according to a braking profile. In some embodiments, the braking profile may include a plurality of segments (e.g., including a first segment associated with a first deceleration rate, a second segment which includes a second deceleration rate less than the first deceleration rate, and a third segment in which a level of braking is decreased as a target stopping location is approached, as determined based on the analysis of the plurality of images.

[0321] FIG. 32A illustrates vehicle 200 traveling on a roadway 3210 that includes an object 3220. System 100 may identify a trigger (e.g., object 3220) for stopping vehicle 200 in respect to detection of the trigger. For example, system 100 may first slow vehicle 200 after the trigger is detected and bring vehicle 200 to a stop before reaching object 3220.

[0322] Processing unit 110 may be configured to determine a trigger for stopping vehicle 200 based on a plurality of images acquired by one or more of image capture devices 122-126. Processing unit 110 may receive the images from image capture devices 122-126 via data interface 128. According to some embodiments, the trigger may include object 3220. For example, object 3220 may be another vehicle on roadway 2810. According to some embodiments, processing unit 110 may identify a trigger if object 3220 is within a certain vicinity of vehicle 200. For example, identification of the trigger may depend on a relative distance between object 3220 and vehicle 200. Further, this identification may optionally consider the relative velocity and/or relative acceleration between object 3220 and vehicle 200.

[0323] Additionally or alternatively, the trigger may include a traffic sign, such as traffic sign 3240 shown in FIG. 32A. For example, traffic sign 3240 may be a stop sign and/or a speed limit sign. According to some embodiments, processing unit 110 may process the plurality of images to determine the content of traffic sign 3240 as part of determining whether a trigger exists. For example, processing unit 110 may determine that a trigger is present based on a speed limit specified by traffic sign 3240 if vehicle 200 is exceeding the speed limit. As another example, processing unit 110 may determine a trigger exists based on a traffic light 2830 if the yellow and/or red light is on.

[0324] Based on the plurality of images, processing unit 110 may determine a target stopping location 3250 at which vehicle 200 will stop once a braking profile is

IPR2025-00943
Tesla EX1002 Page 1719

executed. Reaching target stopping location 3250 may be accompanied by a release in braking pressure of vehicle 200. Target stopping location 3250 may be based on the location of traffic light 3230 and/or traffic sign 3240 such that vehicle 200 complies with traffic light 3230 and/or traffic sign 3240. Alternatively, target stopping location 3250 may be based on object 3220 such that vehicle 200 stops before reaching object 3220. Accordingly, target stopping location 3250 may be based on the location of object 3220 to avoid a collision between vehicle 200 and object 3220. In some embodiments, target stopping location 3250 may be updated based on new information obtained from the plurality of images, such as newly identified targets, newly visible obstacles, newly visible signs, etc.

[0325] FIG. 32B illustrates an exemplary braking profile, consistent with disclosed embodiments. In the example shown in FIG. 32B, the braking profile includes four segments. However, the number of segments is exemplary and a braking profile consistent with the disclosed embodiments may include any appropriate number of segments (e.g., 2, 3, 4, 5, 6, etc., segments).

[0326] In segment 1, system 100 may prepare for a possible braking scenario. For example, in this zone, system 100 may receive information (e.g., map data and/or image data) about an upcoming intersection (e.g., an intersection with traffic lights). System 100 may gradually reduce the speed of vehicle 200 (e.g., the speed of vehicle 200 may be gradually reduced to about 70-50 km/hr, depending on the characteristics of the intersection). System 100 may cause the speed adjustment, for example, approximately 100-200 meters before the junction at average deceleration around 0.2-0.4 m/sec^2 (with a maximum of about 0.5 m/sec^2). By decreasing the speed of vehicle 200 in segment 1, system 100 may provide the driver and/or passengers of vehicle 200 with confidence that system 100 is aware of the approaching intersection.

[0327] In segment 2, system 100 may apply strong braking. In this segment, system 100 may receive information based on the analysis of image data that, for example, a traffic light is red and/or there is a stopped vehicle at the intersection in the line in which vehicle 200 is traveling. System 100 may cause vehicle 200 to experience significant speed reduction in this zone, while the distance to the stop line or stopped vehicle is still large. Doing so may provide the driver and/or passengers of vehicle 200 with a comfortable feeling that system 100 will have more than enough space to complete the braking maneuver. This zone may be approximately 30-100 meters away from the stop line, and the average deceleration may be about 1.5-2.5 m/sec^2 (with a maximum of about 3.5 m/sec^2).

[0328] In segment 3, system 100 may make a moderate braking adjustment. For example, in this segment, system 100 may adjust the speed of vehicle 200 according to the remaining distance to the stop line or stopped vehicle and based on the current speed and deceleration of vehicle 200. By making this adjustment, system 100 may provide the driver and/or passengers of vehicle 200 with an indication that system 100 is releasing most of the braking power while the speed is still being reduced but slower than before. This zone may be approximately 5-30 meters away from the stop line or stopped vehicle, and the average deceleration may be about 0.5-1.5 m/sec^2 (with a maximum of about 2 m/sec^2).

[0329] In segment 4, system 100 may make a small braking adjustment. In this segment, system 100 may close the remaining distance to the stop line or stopped vehicle at a very low speed gradually bringing it to zero. In this zone, the driver and/or passengers may feel that vehicle 200 is slowly sliding into its position (e.g., at the stop line or behind another vehicle). This zone may constitute approximately the last 5-7 meters to the stop line or stopped vehicle, and the average deceleration may be about 0.3-0.5 m/sec^2 (with a maximum of about 1 m/sec^2).

[0330] FIG. 33 is an exemplary block diagram of memory 140 and/or 150, which may store instructions for performing one or more operations consistent with disclosed embodiments. As illustrated in FIG. 33, memory 140 may store one or more modules for performing the trigger detection and responses described herein. For example, memory 140 may store a trigger identification module 3300 and an action response module 3310.

[0331] Trigger identification module 3300 may store instructions which, when executed by processing unit 110, may detect the presence and/or existence of a trigger for stopping vehicle 200. For example, trigger identification module 3300 may process the plurality of images received from at least one image capture device 122-124 to detect the presence of a trigger. As discussed above, this may include identifying object 3220, traffic light 3230, and/or traffic sign 3240.

[0332] Action response module 3310 may store instructions which, when executed by processing unit 110, may respond to the presence of a trigger. For example, action response module 3310 may execute control to decelerate and/or stop vehicle 200. For example, action response module 3310 may execute a braking profile to decelerate and/or stop vehicle 200. The action response module 3310 may determining to execute one of a plurality of braking profiles based on a number of braking profiles based on, for example, environmental factors, the type of trigger detected, and the velocity of vehicle 200. To navigate vehicle 200 to a stop based on the detected trigger, processing unit 110 may transmit electronic signals to one or more of throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240. For example, processing unit 110 may transmit electronic signals that cause system 100 to physically depress the brake by a predetermined amount.

[0333] Fig. 34 illustrates a process 3400 for navigating vehicle 200, consistent with disclosed embodiments. Process 3000 may identify a trigger (e.g., an object, ve-

IPR2025-00943
Tesla EX1002 Page 1720

hicle, traffic light, traffic sign, etc.) for stopping vehicle 200 and execute a braking profile in response to the detection of the trigger. Process 3400 may detect various types of triggers determine which of a plurality of braking profiles to execute.

[0334] The braking profile may be based on the type of a detected trigger and/or other environmental conditions related to the detected trigger (e.g., a distance to an object). For example, in some embodiments, the braking profile may include a plurality of segments (e.g., three segments). In a three segment braking profile, a first segment of the braking profile may be associated with a first deceleration rate, a second segment of the braking profile may include a second deceleration rate less than the first deceleration rate, and a third segment of the braking profile may include causing a level of braking to be decreased as a target stopping location is approached, as determined based on an analysis of a plurality of images acquired by one or more of image capture devices 122-126. For example, in some embodiments, the first segment may cause vehicle 200 to progressively increase deceleration, the second segment may cause vehicle 200 to decelerate at a constant rate, and the third segment may cause vehicle 200 to progressively decrease its deceleration rate to zero (or near zero).

[0335] At step 3410, process 3400 may acquire, using at least one of image capture devices 122-126, a plurality of images of an area in the vicinity of vehicle 200. For example, processing unit 110 may receive the plurality of images may through data interface 128. Processing unit 110 may be configured to determine the presence of one or more objects and/or signs located on or near roadway 3210 based on the acquired images. For example, processing unit 110 may be configured to determine the presence of object 3220, traffic light 3230, and/or traffic sign 3240. Processing unit 110 may be further configured to read the content of traffic sign 3240. For example, processing unit 110 may be configured to differentiate between a speed limit sign and a stop sign. Further, processing unit 110 may be configured to process images to determine the speed limit based on traffic sign 3240.

[0336] At step 3420, process 300 may include a trigger for stopping vehicle 200. For example, processing unit 110 may identify a trigger based on the presence of object 3220 within the vicinity of vehicle 200. Additionally or alternatively, processing unit 110 may identify a trigger based on a current signal status of traffic light 3230. According to some embodiments, processing unit 110 may determine that traffic light 3230 constitutes a trigger if it is not green. According to some embodiments, processing unit 110 may determine that traffic sign 3240 constitutes a trigger if it is a stop sign, for example. In some embodiments, traffic sign 3240 may constitute a trigger if it is a speed limit sign and vehicle 200 is exceeding the posted speed limit. Further in some embodiments, processing unit 110 may cause system 100 to provide an audible announcement identifying the trigger and/or indicating that vehicle 200 is braking or about to begin braking.

[0337] After processing unit 110 identifies the trigger, at step 3430, processing unit 110 may cause vehicle 200 to slow down and/or stop according to a braking profile. In some embodiments, the braking profile may comprise a number of profile segments. The different profile segments of a braking profile may be defined by a time period and/or a vehicle speed. For example, processing unit 110 may execute controls of vehicle 200 in accordance with a first profile segment for a predetermined period of time or until the speed of vehicle 200 reaches a target value. The target value may be a predetermined speed or it may be a percentage of the speed that vehicle 200 was traveling prior to execution of the braking profile. To implement the braking profile, processing unit 110 may transmit electronic signals to one or more of throttling system 220, braking system 230, and/or steering system 240 of vehicle 200. Further, one or more actuators may control throttling system 220, braking system 230, and/or steering system 240.

[0338] As discussed above, according to some embodiments, a profile segment of the braking profile may correspond to a deceleration rate. For example, a braking profile may include multiple segments, and the first segment may be associated with a first deceleration rate. The braking profile may further include a second segment associated with a second deceleration rate. The second deceleration rate may be less than first deceleration rate.

[0339] Further, as discussed above, according to some embodiments, braking profile may include a third profile segment in which a level of braking is decreased as a function of the distance between vehicle 200 and underlying factors that give rise to the trigger (e.g., object 3220, traffic light 3230, traffic sign 3240). For example, the braking profile may include decreasing the level of braking as vehicle 200 approaches target stopping location 3250. Further, reaching target stopping location 3250 may be accompanied by a release in braking pressure of vehicle 200.

[0340] The foregoing description has been presented for purposes of illustration. It is not exhaustive and is not limited to the precise forms or embodiments disclosed. Modifications and adaptations will be apparent to those skilled in the art from consideration of the specification and practice of the disclosed embodiments. Additionally, although aspects of the disclosed embodiments are described as being stored in memory, one skilled in the art will appreciate that these aspects can also be stored on other types of computer readable media, such as secondary storage devices, for example, hard disks or CD ROM, or other forms of RAM or ROM, USB media, DVD, Blu-ray, or other optical drive media.

[0341] Computer programs based on the written description and disclosed methods are within the skill of an experienced developer. The various programs or program modules can be created using any of the techniques known to one skilled in the art or can be designed in

IPR2025-00943
Tesla EX1002 Page 1721

connection with existing software. For example, program sections or program modules can be designed in or by means of .Net Framework, Net Compact Framework (and related languages, such as Visual Basic, C, etc.), Java, C++, Objective-C, HTML, HTML/AJAX combinations, XML, or HTML with included Java applets.

Claims

1. A driver assist navigation system (100) for a primary vehicle (200a), the system comprising:

at least one image capture device (122, 124, 126) configured to acquire a plurality of images of an area in a vicinity of the primary vehicle;
a data interface (128); and
at least one processing device (110) configured to:

receive (1810) the plurality of images via the data interface (128);
determine (1820), from at least some of the plurality of images, a first lane constraint (810) on a first side (802) of the primary vehicle (200a);
determine (1830), from at least some of the plurality of images, a second lane constraint (820) on a second side (812) of the primary vehicle (200a) opposite to the first side (802) of the primary vehicle, wherein the first and second lane constraints define a lane within which the primary vehicle (200a) travels; and
send control signals to the control systems (220, 230, 240) of said primary vehicle for causing (1840) the primary vehicle (200a) to travel within the first and second lane constraints (810, 820);

wherein:

the first lane constraint includes a line marked on a road surface, an edge of a road surface, markers on a road surface, or is estimated based on a determined midpoint of a road surface width;
the second lane constraint includes a line marked on a road surface, an edge of a road surface, markers on a road surface, or is estimated based on a determined midpoint of a road surface width; and
characterized in that
the at least one processing device (110) is further configured to:

locate (1850) in the plurality of images a leading vehicle (200b);
determine (1860), based on the plurality of images, at least one action taken by the leading vehicle (200b); and
send control signals to the control systems (220, 230, 240) for causing (1870) the primary vehicle (200a) to mimic the at least one action of the leading vehicle (200b);

wherein the at least one action taken by the leading vehicle (200b) includes performing a maneuver such that a first offset distance (c_3) on a side of the leading vehicle (200b) adjacent to the first lane constraint (810) is different from a second offset distance (c_4) on a side of the leading vehicle (200b) adjacent to the second lane constraint (820), and
wherein mimicking the at least one action includes the at least one processing device causing the primary vehicle to perform a maneuver such that a third offset distance (c_1) on a side of the primary vehicle (200a) adjacent to the first lane constraint (810) is different from a fourth offset distance (c_2) on a side of the primary vehicle (200a) adjacent to the second lane constraint (820), wherein the third offset distance (c_1) equals to the first offset distance (c_3) and the fourth offset distance (c_2) equals to the second offset distance (c_4).

2. The system of claim 1, wherein the at least one action taken by the leading vehicle (200b) further includes a turn, and mimicking the at least one action respectively includes the at least one processing device (110) causing the primary vehicle (200a) to execute a similar turn beginning at a location at which the leading vehicle (200b) initiated its turn.

3. The system of claim 1, wherein the at least one action taken by the leading vehicle further includes an acceleration and wherein mimicking the at least one action respectively includes the at least one processing device causing the primary vehicle to accelerate.

4. The system of claim 1, wherein the at least one action taken by the leading vehicle further includes braking wherein mimicking the at least one action respectively includes the at least one processing device causing the primary vehicle to brake.

5. The system of claim 1, wherein the at least one processing device (110) is further configured to:

determine (1925), based on the plurality of images or based on map data and position infor-

IPR2025-00943
Tesla EX1002 Page 1722

mation, whether the leading vehicle (200b) is turning at an intersection; and
if the leading vehicle (200b) is determined to be turning at an intersection, decline (1930) to cause the primary vehicle to perform a turn similar to the turn of the leading vehicle (200b).

6. The system of claim 1, wherein the at least one processing device (110) is further configured to:

determine (2025), based on the plurality of images or based on map data and position information, whether the leading vehicle (200b) is turning at an intersection;
determine (2035) whether the leading vehicle (200b) is changing lanes within the intersection;
mimic the turn of the leading vehicle (200b) by causing the primary vehicle (200a) to execute a similar turn beginning at a location at which the leading vehicle (200b) initiated its turn, if the leading vehicle (200b) is determined to be turning at an intersection and not changing lanes within the intersection; and
decline (2050) to cause the primary vehicle to perform a turn similar to the turn of the leading vehicle (200b) if the leading vehicle (200b) is determined to be turning at an intersection, but changing lanes within the intersection.

7. The system of claim 1, wherein:

the leading vehicle is a first leading vehicle;
at least one processing device (110) is further configured to: locate (2130) a second leading vehicle in the plurality of images; and
cause the primary vehicle (200a) to mimic the at least one action of the first leading vehicle if, based on the plurality of images, the first leading vehicle is determined to be continuing on a path having less of a turning radius than a path followed by the second leading vehicle.

8. The system of claim 1, wherein the at least one processing device (110) is further configured to:

determine (2230) position information for the leading vehicle (200b) and compare (2240) the position information to predetermined map data;
create (2250) a navigation history for the leading vehicle (200b) by tracking movements of the leading vehicle (200b) relative to the predetermined map data; and
determine (2260) whether to cause the primary vehicle (200a) to mimic the at least one action of the leading vehicle (200b) based on the navigation history of the leading vehicle (200b).

9. The system of claim 8, wherein the at least one processing device (110) is configured to:

cause the primary vehicle (200a) to perform a turn similar to a turn of the leading vehicle (200b), even where no turn is detected based on the plurality of images, when the navigation history indicates that the leading vehicle (200b) has successfully navigated at least one prior turn.

10. A primary vehicle for following a leading vehicle, the primary vehicle comprising the system of any preceding claim.

11. A method for navigating a primary vehicle (200a), the method comprising:

acquiring (1810), using at least one image capture device (122, 124, 126), a plurality of images of an area in a vicinity of the primary vehicle (200a);
determining (1820), from at least some of the plurality of images, a first lane constraint (810) on a first side (802) of the primary vehicle (200a);
determining (1830), from at least some of the plurality of images, a second lane constraint (820) on a second side (812) of the primary vehicle (200a) opposite to the first side (802) of the primary vehicle (200a), wherein the first and second lane constraints (810, 820) define a lane within which the primary vehicle (200a) travels; and
sending control signals to the control systems (220, 230, 240) of said primary vehicle for causing (1840) the primary vehicle (200a) to travel within the first and second lane constraints (810, 820); wherein:
the first lane constraint includes a line marked on a road surface, an edge of a road surface, markers on a road surface, or is estimated based on a determined midpoint of a road surface width;
the second lane constraint includes a line marked on a road surface, an edge of a road surface, markers on a road surface, or is estimated based on a determined midpoint of a road surface width, and
characterized in that the method further comprises:

locating (1850) in the plurality of images a leading vehicle (200b);
determining (1860), based on the plurality of images, at least one action taken by the leading vehicle (200b); and
sending control signals to the control systems (220, 230, 240) for causing (1870) the primary vehicle (200a) to mimic the at least

IPR2025-00943
Tesla EX1002 Page 1723

one action of the leading vehicle (200b);

wherein the at least one action taken by the leading vehicle (200b) includes performing a maneuver such that a first offset distance (c_3) on a side of the leading vehicle (200b) adjacent to the first lane constraint (810) is different from a second offset distance (c_4) on a side of the leading vehicle (200b) adjacent to the second lane constraint (820), and wherein mimicking the at least one action includes the at least one processing device causing the primary vehicle to perform a maneuver such that a third offset distance (c_1) on a side of the primary vehicle (200a) adjacent to the first lane constraint (810) is different from a fourth offset distance (c_2) on a side of the primary vehicle (200a) adjacent to the second lane constraint (820), wherein the third offset distance (c_1) equals to the first offset distance (c_3) and the fourth offset distance (c_2) equals to the second offset distance (c_4).

12. The method of claim 11, further comprising:

determining (1925), based on the plurality of images or based on map data and position information, whether the leading vehicle (200b) is turning at an intersection; and
if the leading vehicle (200b) is determined to be turning at an intersection, declining (1930) to mimic the turn of the leading vehicle (200b).

13. The method of claim 11, further comprising:

determining (2025), based on the plurality of images or based on map data and position information, whether the leading vehicle (200b) is turning at an intersection;
determining (2035) whether the leading vehicle (200b) is changing lanes within the intersection;
mimicking the turn of the leading vehicle (200b) by causing the primary vehicle (200a) to execute a similar turn beginning at a location at which the leading vehicle (200b) initiated its turn, if the leading vehicle (200b) is determined to be turning at an intersection and not changing lanes within the intersection; and
declining (2050) to cause the primary vehicle to perform a turn similar to the turn of the leading vehicle (200b) if the leading vehicle (200b) is determined to be turning at an intersection, but changing lanes within the intersection.

14. The method of claim 11, wherein the leading vehicle is a first leading vehicle, the method further comprising:

locating (2130) a second leading vehicle in the plurality of images; and
causing the primary vehicle (200a) to mimic the at least one action of first leading vehicle if, based on the plurality of images, the first leading vehicle is determined to be continuing on a path having less of a turning radius than a path followed by the second leading vehicle.

15. The method of claim 11, further comprising to:

determining (2230) position information for the leading vehicle (200b) and compare (2240) the position information to predetermined map data;
creating (2250) a navigation history for the leading vehicle (200b) by tracking movements of the leading vehicle (200b) relative to the predetermined map data; and
determining (2260) whether to cause the primary vehicle (200a) to mimic the at least one action of the leading vehicle (200b) based on the navigation history of the leading vehicle (200b); and optionally causing the primary vehicle (200a) to mimic a turn of the leading vehicle (200b), even where no turn is detected based on the plurality of images, when the navigation history indicates that the leading vehicle (200b) has successfully navigated at least one prior turn.

Patentansprüche

1. Fahrerassistenz-Navigationssystem (100) für ein Primärfahrzeug (200a), wobei das System Folgendes umfasst:

mindestens eine Bilderfassungsvorrichtung (122, 124, 126), die dazu konfiguriert ist, eine Vielzahl von Bildern eines Bereiches in einer Nähe des Primärfahrzeugs zu erfassen;
eine Datenschnittstelle (128); und
mindestens eine Verarbeitungsvorrichtung (110), die dazu konfiguriert ist:

die Vielzahl von Bildern über die Datenschnittstelle (128) zu empfangen (1810);
von zumindest einigen der Vielzahl von Bildern eine erste Fahrspurbeschränkung (810) an einer ersten Seite (802) des Primärfahrzeugs (200a) festzustellen (1820);
von zumindest einigen der Vielzahl von Bildern eine zweite Fahrspurbeschränkung (820) an einer zweiten Seite (812) des Primärfahrzeugs (200a), die der ersten Seite

IPR2025-00943
Tesla EX1002 Page 1724

(802) des Primärfahrzeugs entgegengesetzt ist, festzustellen (1830), wobei die erste und zweite Fahrspurbeschränkung eine Fahrspur definieren, innerhalb welcher das Primärfahrzeug (200a) fährt; und

Steuersignale an die Steuersysteme (220, 230, 240) des Primärfahrzeugs zu senden, um zu veranlassen (1840), dass das Primärfahrzeug (200a) innerhalb der ersten und zweiten Fahrspurbeschränkung (810, 820) fährt;

wobei

die erste Fahrspurbeschränkung eine auf einer Fahrbahnoberfläche markierte Linie, einen Rand einer Fahrbahnoberfläche, Markierungen auf einer Fahrbahnoberfläche umfasst, oder basierend auf einem festgestellten Mittelpunkt einer Fahrbahnoberflächenbreite geschätzt wird;

die zweite Fahrspurbeschränkung eine auf einer Fahrbahnoberfläche markierte Linie, einen Rand einer Fahrbahnoberfläche, Markierungen auf einer Fahrbahnoberfläche umfasst, oder basierend auf einem festgestellten Mittelpunkt einer Fahrbahnoberflächenbreite geschätzt wird; und

dadurch gekennzeichnet, dass mindestens eine Verarbeitungsvorrichtung (110) ferner dazu konfiguriert ist:

in der Vielzahl von Bilder ein führendes Fahrzeug (200b) ausfindig zu machen (1850);

basierend auf der Vielzahl von Bildern, mindestens eine durch das führende Fahrzeug (200b) ausgeführte Handlung festzustellen (1860); und

Steuersignale an die Steuersysteme (220, 230, 240) des Primärfahrzeugs zu senden, um zu veranlassen (1870), dass das Primärfahrzeug (200a) mindestens eine Handlung des führenden Fahrzeugs (200b) nachahmt;

wobei die mindestens eine durch das führende Fahrzeug (200b) ausgeführte Handlung das Ausführen eines Manövers umfasst, so dass sich ein erster Versatzabstand (c_3) an einer Seite des führenden Fahrzeugs (200b) angrenzend an die erste Fahrspurbeschränkung (810) von einem zweiten Versatzabstand (c_4) an einer Seite des führenden Fahrzeugs (200b) angrenzend an die zweite Fahrspurbeschränkung (820) unterscheidet, und

wobei das Nachahmen der mindestens einen Handlung umfasst, dass die mindestens eine Verarbeitungsvorrichtung veranlasst, dass das Primärfahrzeug ein Manöver ausführt, so dass sich ein dritter Versatzabstand (c_1) an einer Seite des Primärfahrzeugs (200a) angrenzend an die erste Fahrspurbeschränkung (810) von einem vierten Versatzabstand (c_2) an einer Seite des Primärfahrzeugs (200a) angrenzend an die zweite Fahrspurbeschränkung (820) unterscheidet,

wobei der dritte Versatzabstand (c_1) gleich dem ersten Versatzabstand (c_3) ist und der vierte Versatzabstand (c_2) gleich dem zweiten Versatzabstand (c_4) ist.

2. System nach Anspruch 1, wobei die mindestens eine durch das führende Fahrzeug (200b) ausgeführte Handlung ferner ein Abbiegen umfasst, und das Nachahmen der mindestens einen Handlung jeweils umfasst, dass die mindestens eine Verarbeitungsvorrichtung (110) veranlasst, dass das Primärfahrzeug (200a) ein ähnliches Abbiegen beginnend an einer Stelle ausführt, an welcher das führende Fahrzeug (200b) sein Abbiegen begonnen hat.

3. System nach Anspruch 1, wobei die mindestens eine durch das führende Fahrzeug ausgeführte Handlung ferner eine Beschleunigung umfasst und wobei das Nachahmen der mindestens einen Handlung jeweils umfasst, dass die mindestens eine Verarbeitungsvorrichtung veranlasst, dass das Primärfahrzeug beschleunigt.

4. System nach Anspruch 1, wobei die mindestens eine durch das führende Fahrzeug ausgeführte Handlung ferner ein Bremsen umfasst, wobei das Nachahmen der mindestens einen Handlung jeweils umfasst, dass die mindestens eine Verarbeitungsvorrichtung veranlasst, dass das Primärfahrzeug bremst.

5. System nach Anspruch 1, wobei die mindestens eine Verarbeitungsvorrichtung (110) ferner dazu konfiguriert ist,

basierend auf der Vielzahl von Bildern oder basierend auf Kartendaten und Postitionsinformation festzustellen (1925), ob das führende Fahrzeug (200b) an einer Kreuzung abbiegt; und,

wenn festgestellt wird, dass das führende Fahrzeug (200b) an einer Kreuzung abbiegt, das Veranlassen, dass das Primärfahrzeug ein ähnliches Abbiegen wie das Abbiegen des führenden Fahrzeugs (200b) ausführt, abzulehnen (1930).

6. System nach Anspruch 1, wobei die mindestens eine Verarbeitungsvorrichtung (110) ferner dazu konfiguriert ist,

basierend auf der Vielzahl von Bildern oder basierend auf Kartendaten und Postitionsinformation festzustellen (2025), ob das führende Fahrzeug (200b) an einer Kreuzung abbiegt;

festzustellen (2035), ob das führende Fahrzeug

IPR2025-00943
Tesla EX1002 Page 1725

(200b) Fahrspuren innerhalb der Kreuzung wechselt;
das Abbiegen des führenden Fahrzeugs (200b) durch Veranlassen, dass das Primärfahrzeug (200a) ein ähnliches Abbiegen an einer Stelle ausführt, an welcher das führende Fahrzeug (200b) sein Abbiegen begonnen hat, nachzuahmen, wenn festgestellt wird, dass das führende Fahrzeug (200b) an einer Kreuzung abbiegt und keine Fahrspuren innerhalb der Kreuzung wechselt; und
ein Veranlassen, dass das Primärfahrzeug ein ähnliches Abbiegen wie das Abbiegen des führenden Fahrzeugs (200b) ausführt, abzulehnen (2050), wenn festgestellt wird, dass das führende Fahrzeug (200b) an einer Kreuzung abbiegt, jedoch Fahrspuren innerhalb der Kreuzung wechselt.

7. System nach Anspruch 1, wobei:

das führende Fahrzeug ein erstes führendes Fahrzeug ist;
mindestens eine Verarbeitungsvorrichtung (110) ferner dazu konfiguriert ist: ein zweites führendes Fahrzeug in der Vielzahl von Bildern ausfindig zu machen (2130); und
zu veranlassen, dass das Primärfahrzeug (200a) die mindestens eine Handlung des ersten führenden Fahrzeugs nachahmt, wenn, basierend auf der Vielzahl von Bildern, festgestellt wird, dass das erste führende Fahrzeug auf einer Bahn verbleibt, die einen geringeren Kurvenradius aufweist als eine vom zweiten führenden Fahrzeug gefolgte Bahn.

8. System nach Anspruch 1, wobei die mindestens eine Verarbeitungsvorrichtung (110) ferner dazu konfiguriert ist:

Positionsinformation für das führende Fahrzeug (200b) festzustellen (2230) und die Positionsinformation mit vorbestimmten Kartendaten zu vergleichen (2240);
einen Navigationsverlauf für das führende Fahrzeug (220b) durch Verfolgen von Bewegungen des führenden Fahrzeugs (200b) relativ zu den vorbestimmten Kartendaten zu erzeugen (2250); und
festzustellen (2260), ob veranlasst werden soll, dass das Primärfahrzeug (200a) die mindestens eine Handlung des führenden Fahrzeugs (200b) basierend auf dem Navigationsverlauf des führenden Fahrzeugs (200b) nachahmt.

9. System nach Anspruch 8, wobei die mindestens eine Verarbeitungsvorrichtung (110) dazu konfiguriert ist:

zu veranlassen, dass das Primärfahrzeug (200a) ein Abbiegen ausführt, das einem Abbiegen des führenden Fahrzeugs (200b) ähnlich ist, selbst dort, wo kein Abbiegen basierend auf der Vielzahl von Bildern erkannt wird, wenn der Navigationsverlauf angibt, dass das führende Fahrzeug (200b) mindestens ein vorhergehendes Abbiegen erfolgreich befahren hat.

10. Primärfahrzeug zum Folgen eines führenden Fahrzeugs, wobei das Primärfahrzeug das System nach einem der vorhergehenden Ansprüche umfasst.

11. Verfahren zum Steuern eines Primärfahrzeugs (200a), wobei das Verfahren Folgendes umfasst:

Erfassen (1810), unter Verwendung von mindestens einer Bilderfassungsvorrichtung (122, 124, 126), einer Vielzahl von Bildern eines Bereiches in einer Nähe des Primärfahrzeugs (200a);
Feststellen (1820), von zumindest einigen der Vielzahl von Bildern, einer ersten Fahrspurbeschränkung (810) an einer ersten Seite (802) des Primärfahrzeugs (200a);
Feststellen (1830), von zumindest einigen der Vielzahl von Bildern, einer zweiten Fahrspurbeschränkung (820) an einer zweiten Seite (812) des Primärfahrzeugs (200a), die der ersten Seite (802) des Primärfahrzeugs (200a) entgegengesetzt ist, wobei die erste und zweite Fahrspurbeschränkung (810, 820) eine Fahrspur definieren, innerhalb welcher das Primärfahrzeug (200a) fährt; und
Senden von Steuersignalen an die Steuersysteme (220, 230, 240) des Primärfahrzeugs, um zu veranlassen (1840), dass das Primärfahrzeug (200a) innerhalb der ersten und zweiten Fahrspurbeschränkung (810, 820) fährt;
wobei
die erste Fahrspurbeschränkung eine auf einer Fahrbahnoberfläche markierte Linie, einen Rand einer Fahrbahnoberfläche, Markierungen auf einer Fahrbahnoberfläche umfasst, oder basierend auf einem festgestellten Mittelpunkt einer Fahrbahnoberflächenbreite geschätzt wird;
die zweite Fahrspurbeschränkung eine auf einer Fahrbahnoberfläche markierte Linie, einen Rand einer Fahrbahnoberfläche, Markierungen auf einer Fahrbahnoberfläche umfasst, oder basierend auf einem festgestellten Mittelpunkt einer Fahrbahnoberflächenbreite geschätzt wird; und
dadurch gekennzeichnet, dass das Verfahren ferner Folgendes umfasst:

Ausfindigmachen (1850) eines führenden Fahrzeugs (200b) in der Vielzahl von Bildern;
Feststellen (1860), basierend auf der Viel-

IPR2025-00943
Tesla EX1002 Page 1726

zahl von Bildern, von mindestens einer durch das führende Fahrzeug (200b) ausgeführten Handlung; und

Senden von Steuersignalen an die Steuersysteme (220, 230, 240) des Primärfahrzeugs, um zu veranlassen (1870), dass das Primärfahrzeug (200a) die mindestens eine Handlung des führenden Fahrzeugs (200b) nachahmt;

wobei die mindestens eine durch das führende Fahrzeug (200b) ausgeführte Handlung das Ausführen eines Manövers umfasst, so dass sich ein erster Versatzabstand (c_3) an einer Seite des führenden Fahrzeugs (200b) angrenzend an die erste Fahrspurbeschränkung (810) von einem zweiten Versatzabstand (c_4) an einer Seite des führenden Fahrzeugs (200b) angrenzend an die zweite Fahrspurbeschränkung (820) unterscheidet, und

wobei das Nachahmen der mindestens einen Handlung umfasst, dass die mindestens eine Verarbeitungsvorrichtung veranlasst, dass das Primärfahrzeug ein Manöver ausführt, so dass sich ein dritter Versatzabstand (c_1) an einer Seite des Primärfahrzeugs (200a) angrenzend an die erste Fahrspurbeschränkung (810) von einem vierten Versatzabstand (c_2) an einer Seite des Primärfahrzeugs (200a) angrenzend an die zweite Fahrspurbeschränkung (820) unterscheidet,

wobei der dritte Versatzabstand (c_1) gleich dem ersten Versatzabstand (c_3) ist und der vierte Versatzabstand (c_2) gleich dem zweiten Versatzabstand (c_4) ist.

12. Verfahren nach Anspruch 11, ferner umfassend:

Feststellen (1925), basierend auf der Vielzahl von Bildern oder basierend auf Kartendaten und Postitionsinformation, ob das führende Fahrzeug (200b) an einer Kreuzung abbiegt; und,

wenn festgestellt wird, dass das führende Fahrzeug (200b) an einer Kreuzung abbiegt, Ablehnen (1930), das Abbiegen des führenden Fahrzeugs (200b) nachzuahmen.

13. Verfahren nach Anspruch 11, ferner umfassend:

Feststellen (2025), basierend auf der Vielzahl von Bildern oder basierend auf Kartendaten und Postitionsinformation, ob das führende Fahrzeug (200b) an einer Kreuzung abbiegt;

Feststellen (2035), ob das führende Fahrzeug (200b) Fahrspuren innerhalb der Kreuzung wechselt;

Nachahmen des Abbiegens des führenden Fahrzeugs (200b) durch Veranlassen, dass das Primärfahrzeug (200a) ein ähnliches Abbiegen an einer Stelle ausführt, an welcher das führende Fahrzeug (200b) sein Abbiegen begonnen hat, wenn festgestellt wird, dass das führende Fahrzeug (200b) an einer Kreuzung abbiegt und keine Fahrspuren innerhalb der Kreuzung wechselt; und

Ablehnen (2050) des Veranlassens, dass das Primärfahrzeug ein ähnliches Abbiegen wie das Abbiegen des führenden Fahrzeugs (200b) ausführt, wenn festgestellt wird, dass das führende Fahrzeug (200b) an einer Kreuzung abbiegt, jedoch Fahrspuren innerhalb der Kreuzung wechselt.

14. Verfahren nach Anspruch 11, wobei das führende Fahrzeug ein erstes führendes Fahrzeug ist, wobei das Verfahren ferner Folgendes umfasst:

Ausfindigmachen (2130) eines zweiten führenden Fahrzeugs in der Vielzahl von Bildern; und Veranlassen, dass das Primärfahrzeug (200a) die mindestens ein Handlung des ersten führenden Fahrzeugs nachahmt, wenn, basierend auf der Vielzahl von Bildern, festgestellt wird, dass das erste führende Fahrzeug auf einer Bahn verbleibt, die einen geringeren Kurvenradius aufweist als eine vom zweiten führenden Fahrzeug gefolgte Bahn.

15. Verfahren nach Anspruch 11, ferner umfassend:

Feststellen (2230) von Positionsinformation für das führende Fahrzeug (200b) und Vergleichen (2240) der Positionsinformation mit vorbestimmten Kartendaten;

Erzeugen (2250) eines Navigationsverlaufes für das führende Fahrzeug (200b) durch Verfolgen von Bewegungen des führenden Fahrzeugs (200b) relativ zu den vorbestimmten Kartendaten; und

Feststellen (2260), ob veranlasst werden soll, dass das Primärfahrzeug (200a) die mindestens eine Handlung des führenden Fahrzeugs (200b) basierend auf dem Navigationsverlauf des führenden Fahrzeugs (200b) nachahmt; und

optional Veranlassen, dass das Primärfahrzeug (200a) ein Abbiegen des führenden Fahrzeugs (200b) nachahmt, selbst dort, wo kein Abbiegen basierend auf der Vielzahl von Bildern erkannt wird, wenn der Navigationsverlauf angibt, dass das führende Fahrzeug (200b) mindestens ein vorhergehendes Abbiegen erfolgreich befahren hat.

IPR2025-00943
Tesla EX1002 Page 1727

Revendications

1. Système de navigation d'assistance au conducteur (100) pour un véhicule principal (200a), ledit système comprenant : au moins un dispositif de capture d'image (122, 124, 126) conçu pour acquérir une pluralité d'images d'une zone dans un voisinage du véhicule principal ; une interface de données (128) ; au moins un dispositif de traitement (110) conçu pour :

   recevoir (1810) la pluralité d'images par l'intermédiaire de l'interface de données (128) ;
   déterminer (1820), à partir d'au moins certaines images de la pluralité d'images, une première contrainte de voie (810) sur un premier côté (802) du véhicule principal (200a) ;
   déterminer (1830), à partir d'au moins certaines images de la pluralité d'images, une seconde contrainte de voie (820) sur un second côté (812) du véhicule principal (200a) opposé au premier côté (802) du véhicule principal, lesdites première et seconde contraintes de voie définissant une voie dans laquelle le véhicule principal (200a) circule ;
   et pour envoyer des signaux de commande aux systèmes de commande (220, 230, 240) dudit véhicule principal pour amener (1840) le véhicule principal (200a) à circuler dans les première et seconde contraintes de voie (810, 820) ;
   ladite première contrainte de voie comprenant une ligne marquée sur une surface de chaussée, un bord d'une surface de chaussée, des marqueurs sur une surface de chaussée ou étant estimée sur la base d'un point médian déterminé d'une largeur de surface de chaussée ;
   ladite seconde contrainte de voie comprenant une ligne marquée sur une surface de chaussée, un bord d'une surface de chaussée, des marqueurs sur une surface de chaussée, ou étant estimée sur la base d'un point médian déterminé d'une largeur de surface de chaussée ;
   et **caractérisé en ce que** le au moins un dispositif de traitement (110) est conçu en outre pour: localiser (1850) dans la pluralité d'images un véhicule de tête (200b) ; déterminer (1860), sur la base de la pluralité d'images, au moins une action prise par le véhicule de tête (200b) ;
   et envoyer des signaux de commande aux systèmes de commande (220, 230, 240) dudit véhicule principal pour amener (1870) le véhicule principal (200a) à imiter la au moins une action du véhicule de tête (200b) ;
   ladite au moins une action prise par le véhicule de tête (200b) comprenant la réalisation d'une manoeuvre de sorte qu'une première distance de décalage (c_3) sur un côté du véhicule de tête (200b) adjacent à la première contrainte de voie (810) soit différente d'une seconde distance de décalage (c_4) sur un côté du véhicule de tête (200b) adjacent à la seconde contrainte de voie (820), et ladite imitation de la au moins une action comprenant le au moins un dispositif de traitement amenant le véhicule principal à réaliser une manoeuvre de sorte qu'une troisième distance de décalage (c_1) sur un côté du véhicule principal (200a) adjacent à la première contrainte de voie (810) soit différente d'une quatrième distance de décalage (c_2) sur un côté du véhicule principal (200ba) adjacent à la seconde contrainte de voie (820), ladite troisième distance de décalage (c_1) étant égale à la première distance de décalage (c_3) et ladite quatrième distance de décalage (c_2) étant égale à la seconde distance de décalage (c_4).

2. Système selon la revendication 1, ladite au moins une action prise par le véhicule de tête (200b) comprenant en outre une prise de virage et ladite imitation de la au moins une action respectivement comprenant le au moins un dispositif de traitement (110) amenant le véhicule principal (200a) à exécuter une prise de virage semblable commençant au niveau d'un emplacement au niveau duquel le véhicule de tête (200b) a amorcé sa prise de virage.

3. Système selon la revendication 1, ladite au moins une action prise par le véhicule de tête comprenant en outre une accélération et ladite imitation de la au moins une action respectivement comprenant ledit au moins un dispositif de traitement amenant le véhicule principal à accélérer.

4. Système selon la revendication 1, ladite au moins une action prise par le véhicule de tête comprenant en outre le freinage, ladite imitation de la au moins une action respectivement comprenant le au moins un dispositif de traitement amenant le véhicule principal à freiner.

5. Système selon la revendication 1, ledit au moins un dispositif de traitement (110) étant conçu en outre pour déterminer (1925) sur la base de la pluralité d'images ou sur la base des informations de données cartographiques et de position, si le véhicule de tête (200b) tourne à une intersection ;
   et si le véhicule de tête (200b) est déterminé à tourner à une intersection, pour refuser (1930) d'amener le véhicule principal à effectuer une prise de virage semblable à la prise de virage du véhicule de tête (200b).

6. Système selon la revendication 1, ledit au moins un dispositif de traitement (110) étant conçu en outre pour déterminer (2025) sur la base de la pluralité d'images ou sur la base des informations de don-

IPR2025-00943
Tesla EX1002 Page 1728

nées cartographiques et de position, si le véhicule de tête (200b) tourne à une intersection ;
déterminer (2035) si le véhicule de tête (200b) change de voies dans l'intersection ;
imiter la prise de virage du véhicule de tête (200b) en amenant le véhicule principal (200a) à exécuter une prise de virage semblable commençant au niveau d'un emplacement au niveau duquel le véhicule de tête (200b) a amorcé sa prise de virage, si le véhicule de tête (200 b) est déterminé à tourner à une intersection et à ne pas changer de voie dans l'intersection ;
et refuser (2050) d'amener le véhicule principal à effectuer une prise de virage semblable à la prise de virage du véhicule de tête (200b) si le véhicule de tête (200b) est déterminé à tourner à une intersection mais changeant de voies dans l'intersection.

7. Système selon la revendication 1,
ledit véhicule de tête étant un premier véhicule de tête ;
au moins un dispositif de traitement (110) étant conçu en outre pour : localiser (2130) un second véhicule de tête dans la pluralité d'images ;
et pour amener le véhicule principal (200a) à imiter la au moins une action du premier véhicule de tête si, sur la base de la pluralité d'images, le premier véhicule de tête est déterminé à continuer sur un trajet présentant un rayon de braquage inférieur à celui d'un trajet suivi par le second véhicule de tête.

8. Système selon la revendication 1, ledit au moins un dispositif de traitement (110) étant en outre conçu pour :

déterminer (2230) des informations de position pour le véhicule de tête (200b) et comparer (2240) les informations de position aux données cartographiques prédéfinies ;
créer (2250) un historique de navigation pour le véhicule de tête (200b) en suivant les déplacements du véhicule de tête (200b) par rapport aux données cartographiques prédéfinies ;
et pour déterminer (2260) d'amener ou non le véhicule principal (200a) à imiter la au moins une action du véhicule de tête (200b) sur la base de l'historique de navigation du véhicule de tête (200b).

9. Système selon la revendication 8, ledit au moins un dispositif de traitement (110) étant conçu pour :

amener le véhicule principal (200a) à effectuer une prise de virage semblable à la prise de virage du véhicule de tête (200b), même lorsque aucun virage n'est détecté sur la base de la pluralité d'images, lorsque l'historique de navigation indique que le véhicule de tête (200b) présente une navigation réussie pour au moins une prise de virage précédente.

10. Véhicule principal destiné à suivre un véhicule de tête, ledit véhicule principal comprenant le système selon l'une quelconque des revendications précédentes.

11. Procédé permettant la navigation d'un véhicule principal (200a), ledit procédé comprenant : l'acquisition (1810), à l'aide d'au moins un dispositif de capture d'image (122, 124, 126), d'une pluralité d'images d'une zone dans un voisinage du véhicule principal (200a) ;
la détermination (1820), à partir d'au moins certaines images de la pluralité d'images, d'une première contrainte de voie (810) sur un premier côté (802) du véhicule principal (200a) ;
la détermination (1830), à partir d'au moins certaines images de la pluralité d'images, d'une seconde contrainte de voie (820) sur un second côté (812) du véhicule principal (200a) opposé au premier côté (802) du véhicule principal (200a), lesdites première et seconde contraintes de voie (810, 820) définissant une voie dans laquelle le véhicule principal (200a) circule ;
et l'envoi de signaux de commande aux systèmes de commande (220, 230, 240) dudit véhicule principal pour amener (1840) le véhicule principal (200a) à circuler dans les première et seconde contraintes de voie (810, 820) ;
ladite première contrainte de voie comprenant une ligne marquée sur une surface de chaussée, un bord d'une surface de chaussée, des marqueurs sur une surface de chaussée ou étant estimée sur la base d'un point médian déterminé d'une largeur de surface de chaussée ;
la seconde contrainte de voie comprenant une ligne marquée sur une surface de chaussée, un bord de surface de chaussée, des marqueurs sur une surface de chaussée, ou étant estimée sur la base d'un point médian déterminé de la largeur de surface de chaussée et **caractérisé en ce que** le procédé comprend en outre : la localisation (1850) dans la pluralité d'images d'un véhicule de tête (200b) ;
la détermination (1860), sur la base de la pluralité d'images, d'au moins une action prise par le véhicule de tête (200b) ;
et l'envoi de signaux de commande aux systèmes de commande (220, 230, 240) dudit véhicule principal pour amener (1870) le véhicule principal (200a) à imiter la au moins une action du véhicule de tête (200b) ;
ladite au moins une action prise par le véhicule de tête (200b) comprenant la réalisation d'une manoeuvre de sorte qu'une première distance de décalage (c_3) sur un côté du véhicule de tête (200b) adjacent à la première contrainte de voie (810) est différente

IPR2025-00943
Tesla EX1002 Page 1729

d'une seconde distance de décalage (c_4) sur un côté du véhicule de tête (200b) adjacent à la seconde contrainte de voie (820), et l'imitation de la au moins une action comprenant ledit au moins un dispositif de traitement amenant le véhicule principal à réaliser une manoeuvre de sorte qu'une troisième distance de décalage (c_1) sur un côté du véhicule principal (200a) adjacent à la première contrainte de voie (810) est différente d'une quatrième distance de décalage (c_2) sur un côté du véhicule principal (200a) adjacent à la seconde contrainte de voie (820), ladite troisième distance de décalage (c_1) étant égale à la première distance de décalage (c_3) et la quatrième distance de décalage (c_2) étant égale à la seconde distance de décalage (c_4).

12. Procédé selon la revendication 11, comprenant en outre :

la détermination (1925) sur la base de la pluralité d'images ou sur la base des informations de données cartographiques et de position, pour savoir si le véhicule de tête (200b) tourne à une intersection ;
et si le véhicule de tête (200 b) est déterminé à tourner à une intersection, le refus (1930) d'imiter la prise de virage du véhicule de tête (200b).

13. Procédé selon la revendication 11, comprenant en outre :

la détermination (2025) sur la base de la pluralité d'images ou sur la base des informations de données cartographiques et de position, pour savoir si le véhicule de tête (200b) tourne à une intersection ;
la détermination (2035) pour savoir si le véhicule de tête (200b) change de voies dans l'intersection ; l'imitation de la prise de virage du véhicule de tête (200b) en amenant le véhicule principal (200a) à exécute une prise de virage semblable commençant au niveau d'un emplacement au niveau duquel le véhicule de tête (200b) a amorcé sa prise de virage, si le véhicule de tête (200b) est déterminé à tourner à une intersection et à ne pas changer de voies dans l'intersection ;
et le refus (2050) à amener le véhicule principal à effectuer une prise de virage semblable à la prise de virage du véhicule de tête (200b), si le véhicule de tête (200b) est déterminé à tourner à une intersection, mais changeant de voie dans l'intersection.

14. Procédé selon la revendication 11, ledit véhicule de tête étant un premier véhicule de tête, ledit procédé comprenant en outre :

la localisation (2130) d'un second véhicule de tête dans la pluralité d'images ; et l'incitation du véhicule principal (200a) à imiter la au moins une action du premier véhicule de tête si, sur la base de la pluralité d'images, le premier véhicule de tête est déterminer à continuer sur un trajet présentant un rayon de braquage inférieur à celui du trajet suivi par le second véhicule de tête.

15. Procédé selon la revendication 11, comprenant en outre :

la détermination (2230) des informations de position pour le véhicule de tête (200b) et la comparaison (2240) des informations de position aux données cartographiques prédéfinies ;
la création (2250) d'un historique de navigation pour le véhicule de tête (200b) en suivant les déplacements du véhicule de tête (200b) par rapport aux données cartographiques prédéfinies ;
et la détermination (2260) pour amener ou non le véhicule principal (200a) à imiter la au moins une action du véhicule de tête (200b) sur la base de l'historique de navigation du véhicule de tête (200b) ;
et éventuellement l'incitation du véhicule principal (200a) à imiter une prise de virage du véhicule de tête (200b), même lorsque aucun virage n'est détecté sur la base de la pluralité d'images, lorsque l'historique de navigation indique que le véhicule de tête (200 b) présente une navigation réussie pour au moins une prise de virage précédente.

IPR2025-00943
Tesla EX1002 Page 1730

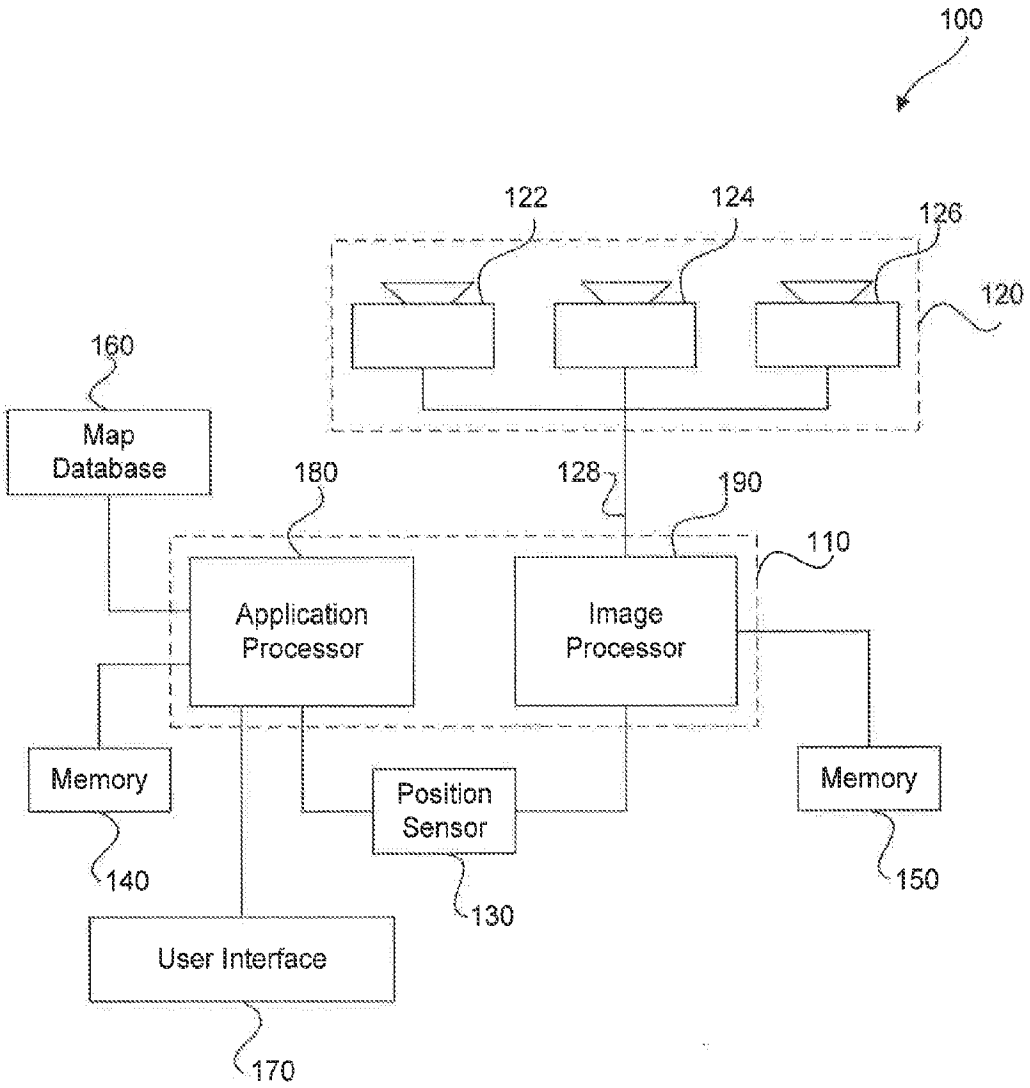


FIG. 1

IPR2025-00943
Tesla EX1002 Page 1731

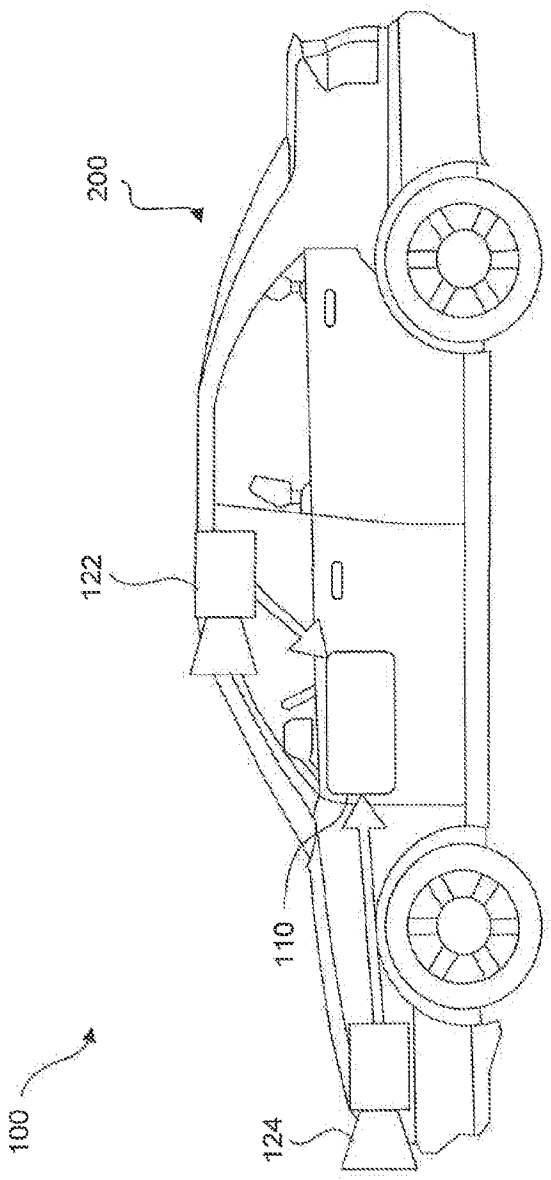


FIG. 2A

IPR2025-00943
Tesla EX1002 Page 1732

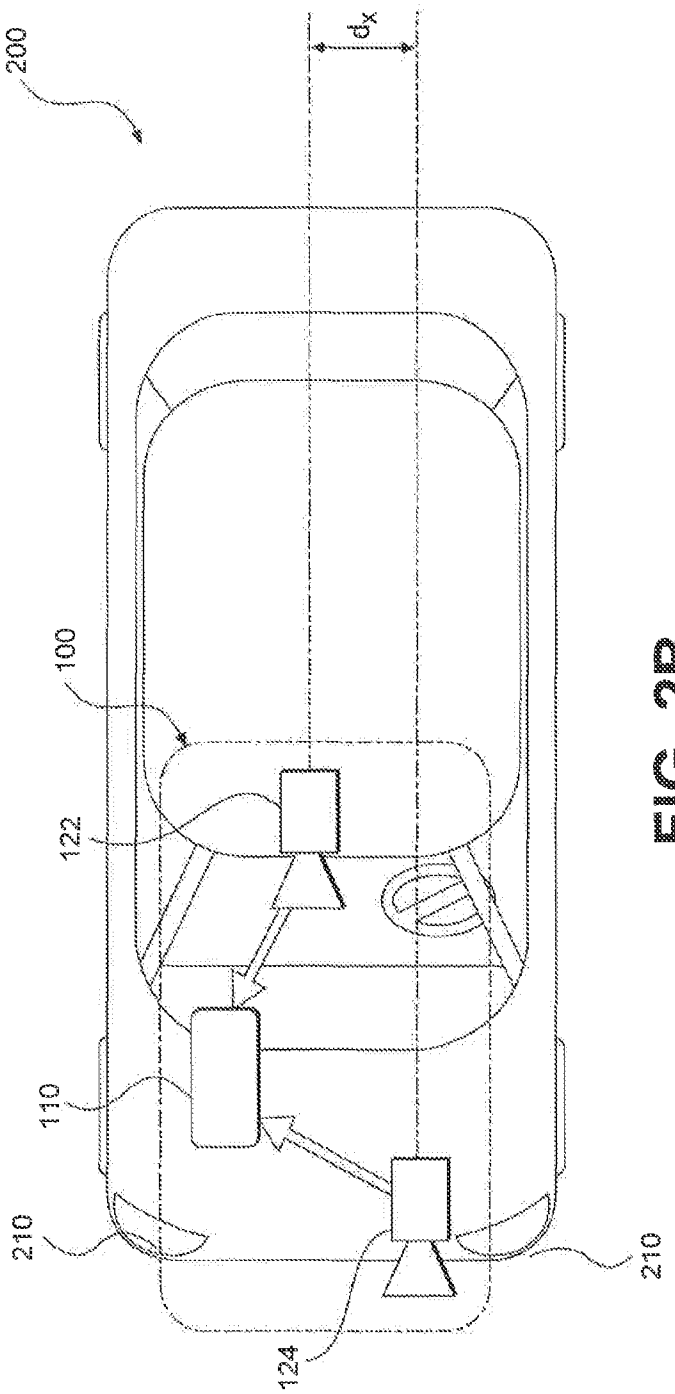


FIG. 2B

IPR2025-00943
Tesla EX1002 Page 1733

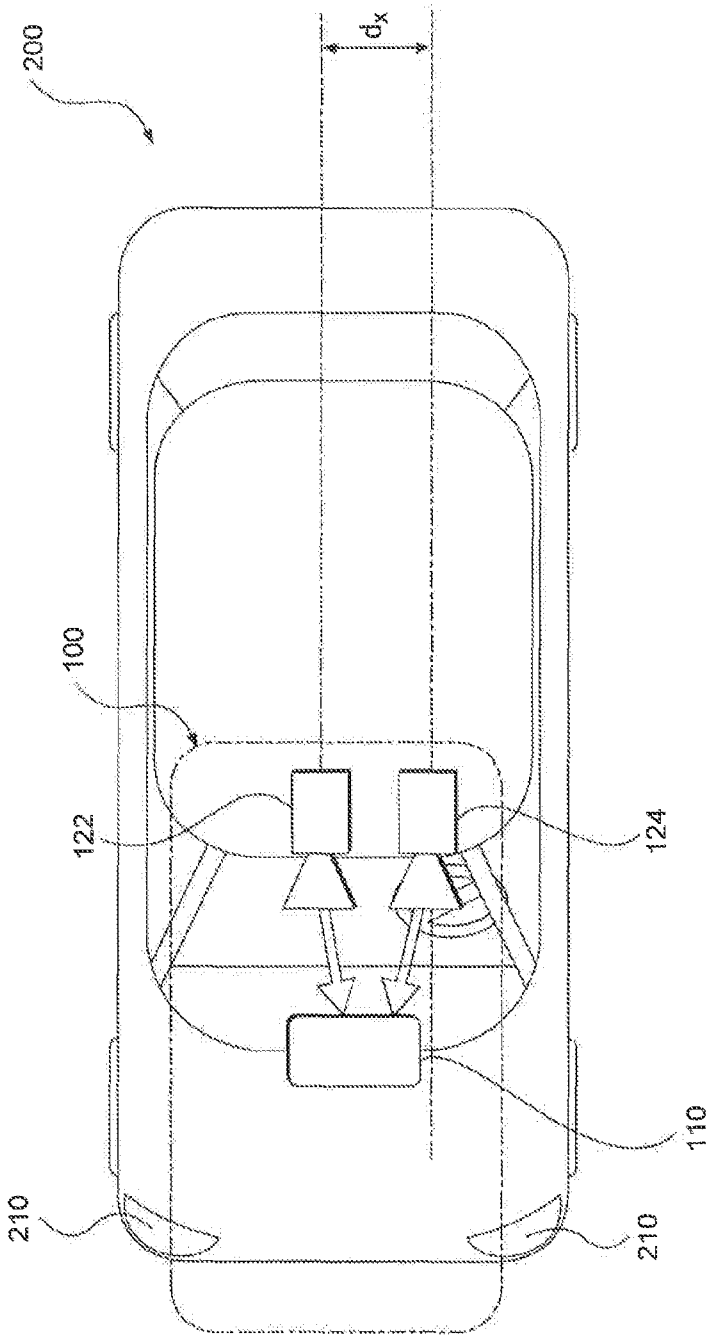


FIG. 2C

IPR2025-00943
Tesla EX1002 Page 1734

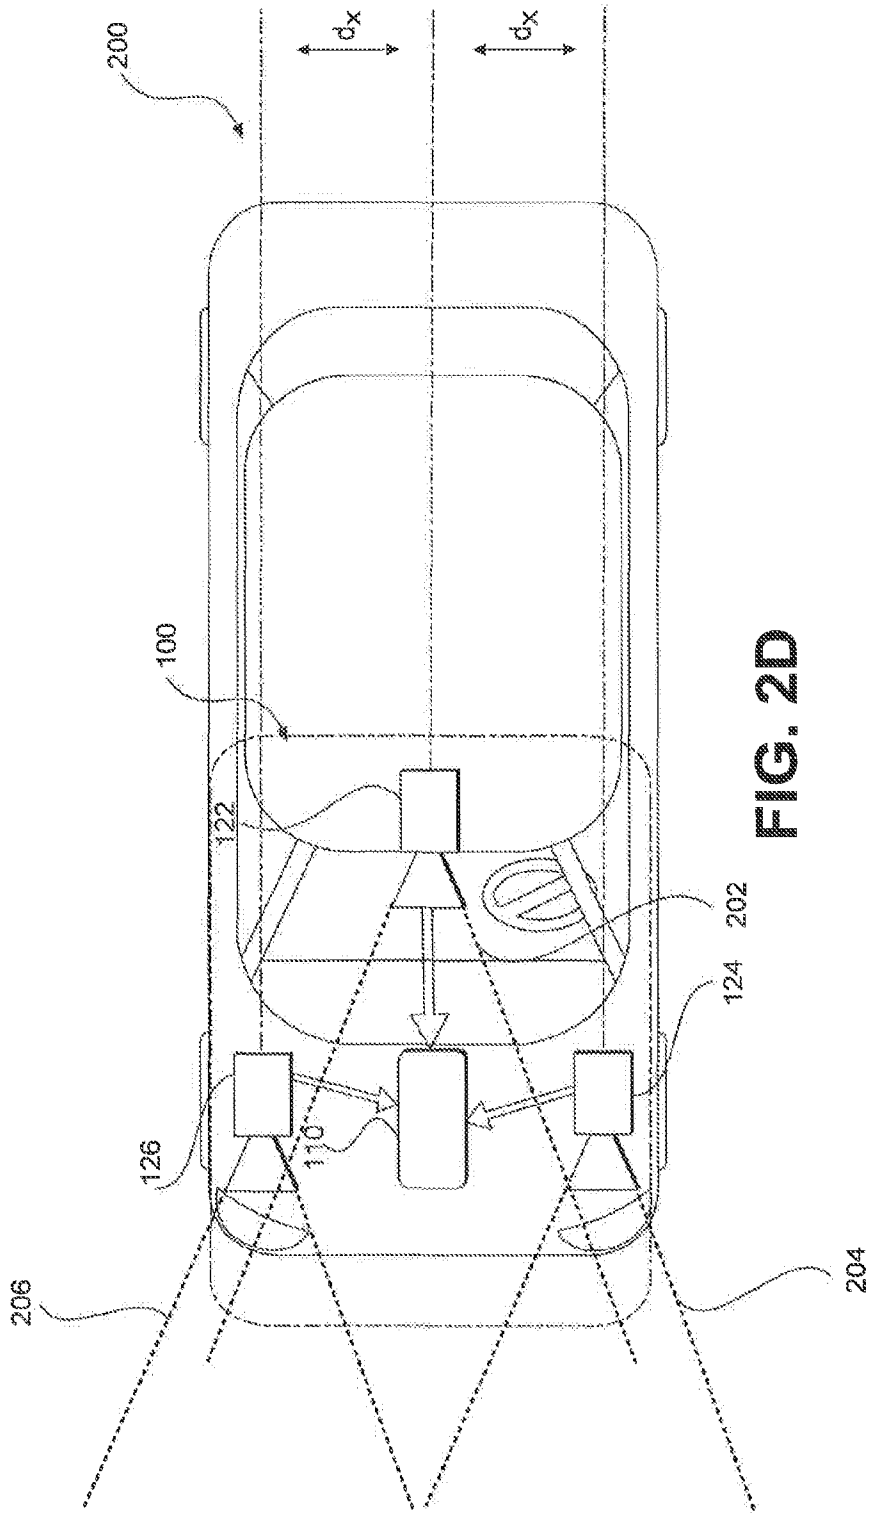


FIG. 2D

IPR2025-00943
Tesla EX1002 Page 1735

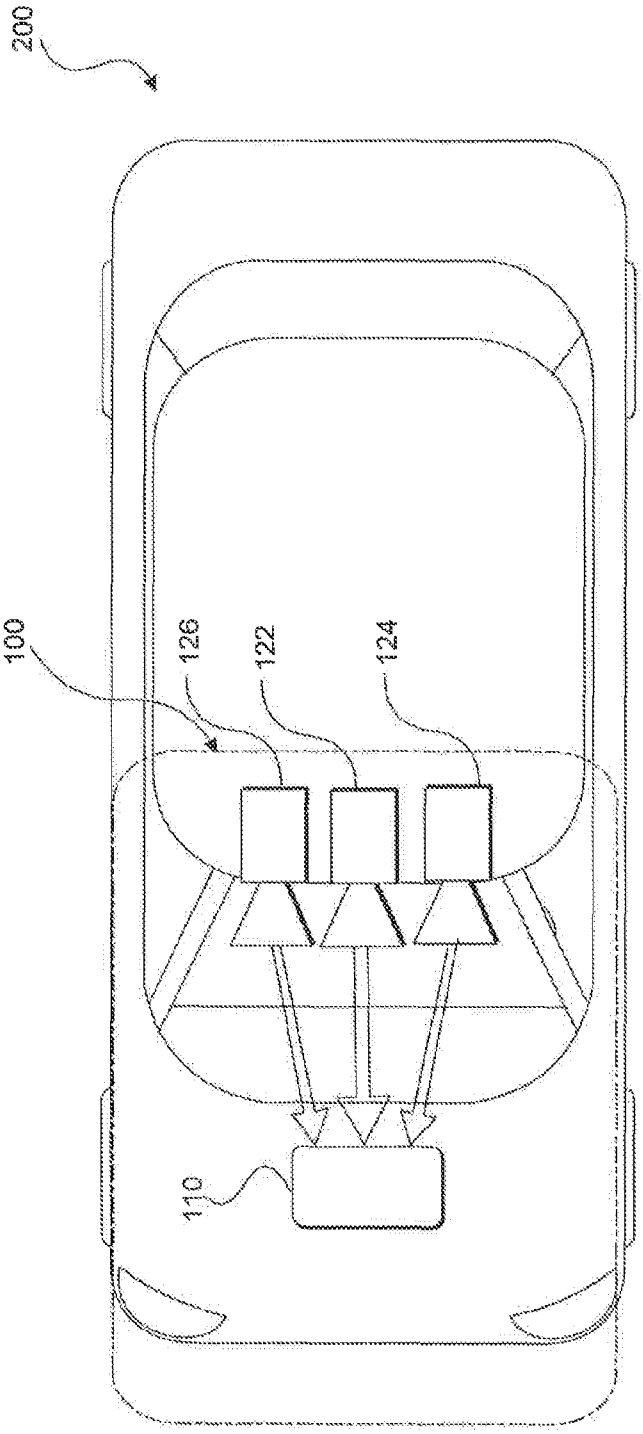


FIG. 2E

IPR2025-00943
Tesla EX1002 Page 1736

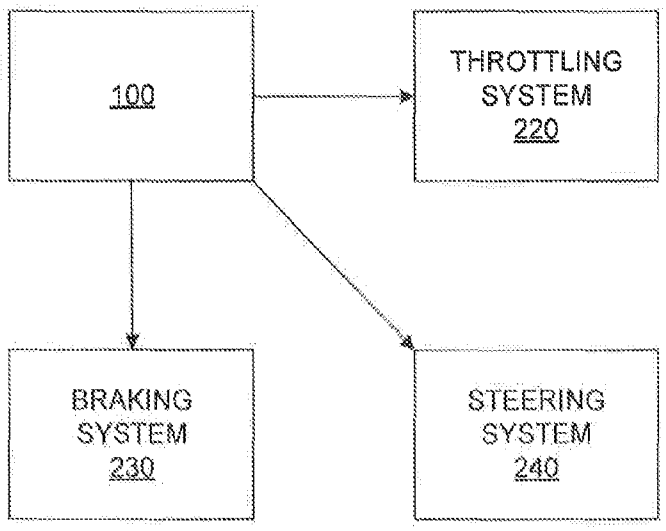


FIG. 2F

IPR2025-00943
Tesla EX1002 Page 1737

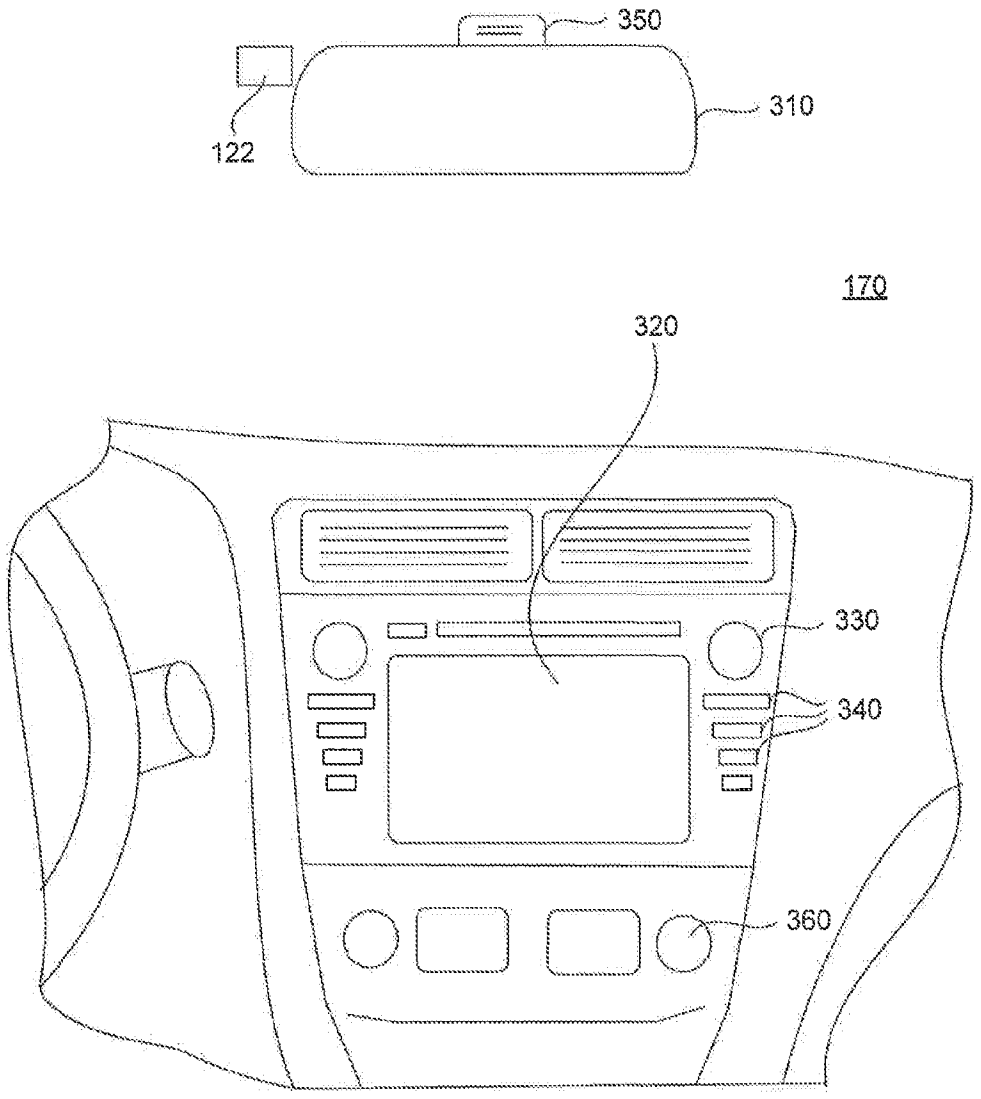


FIG. 3A

IPR2025-00943
Tesla EX1002 Page 1738

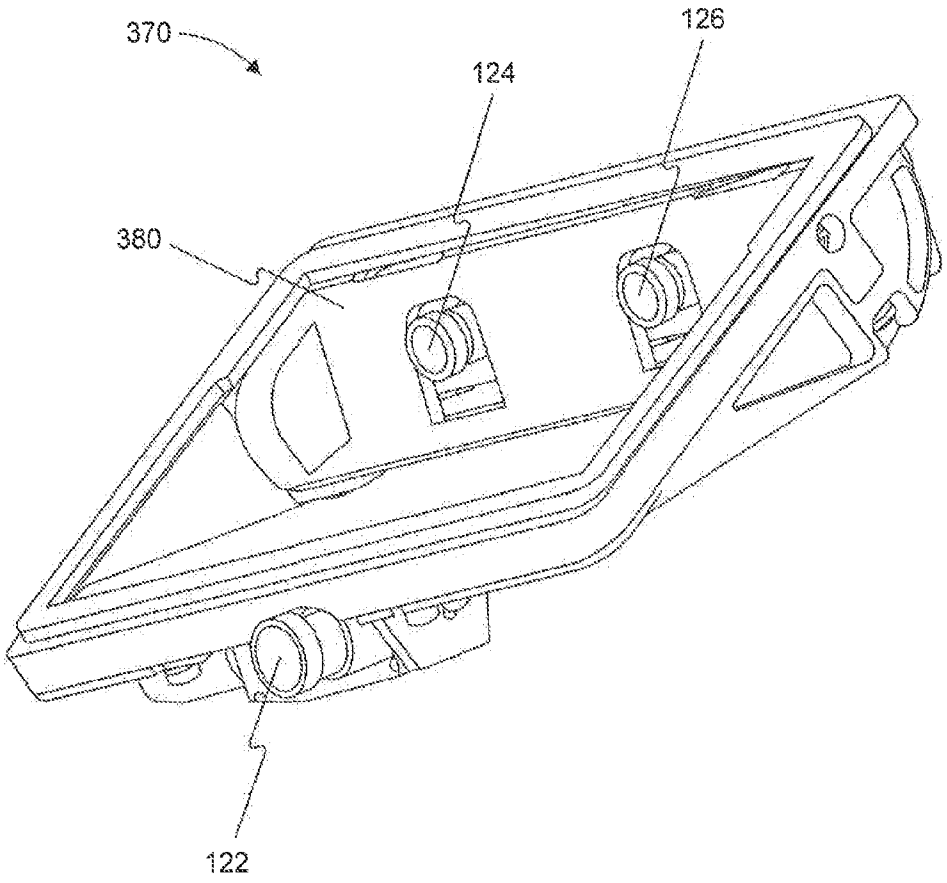


FIG. 3B

IPR2025-00943
Tesla EX1002 Page 1739

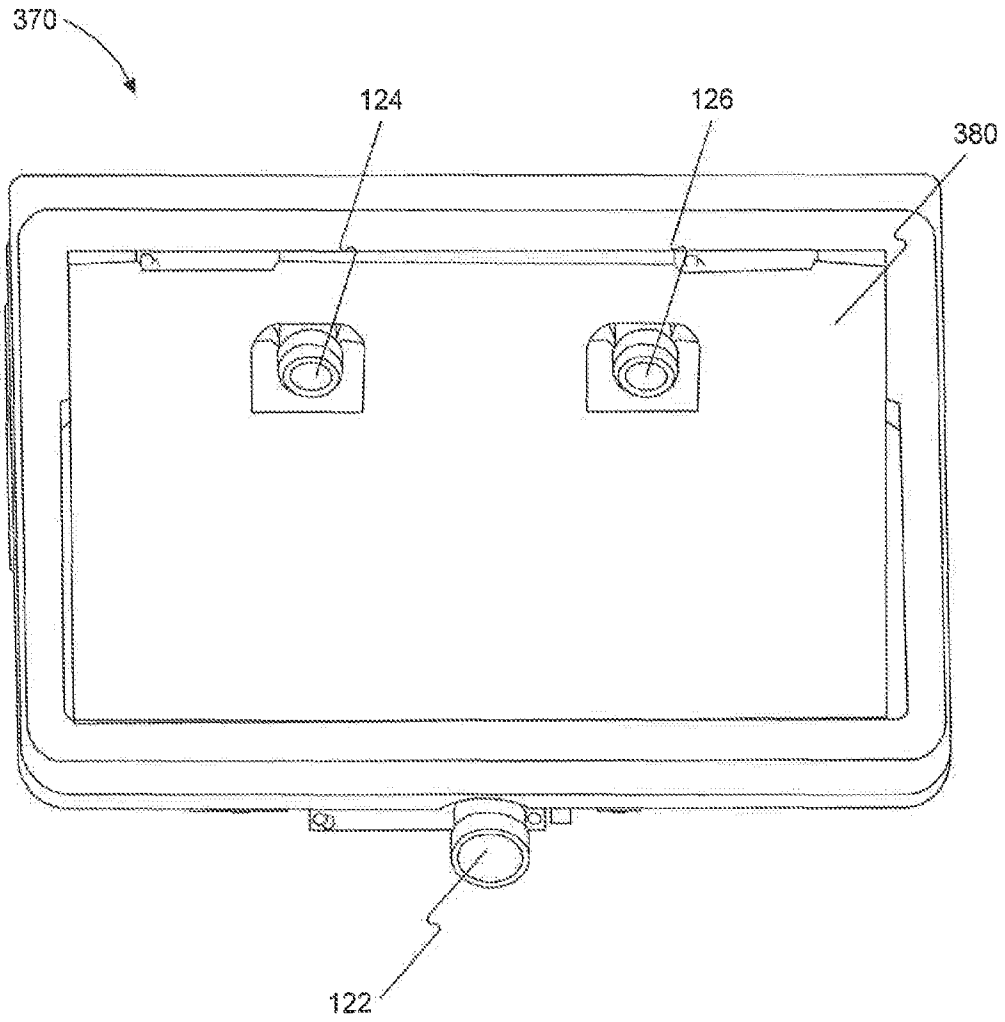


FIG. 3C

IPR2025-00943
Tesla EX1002 Page 1740

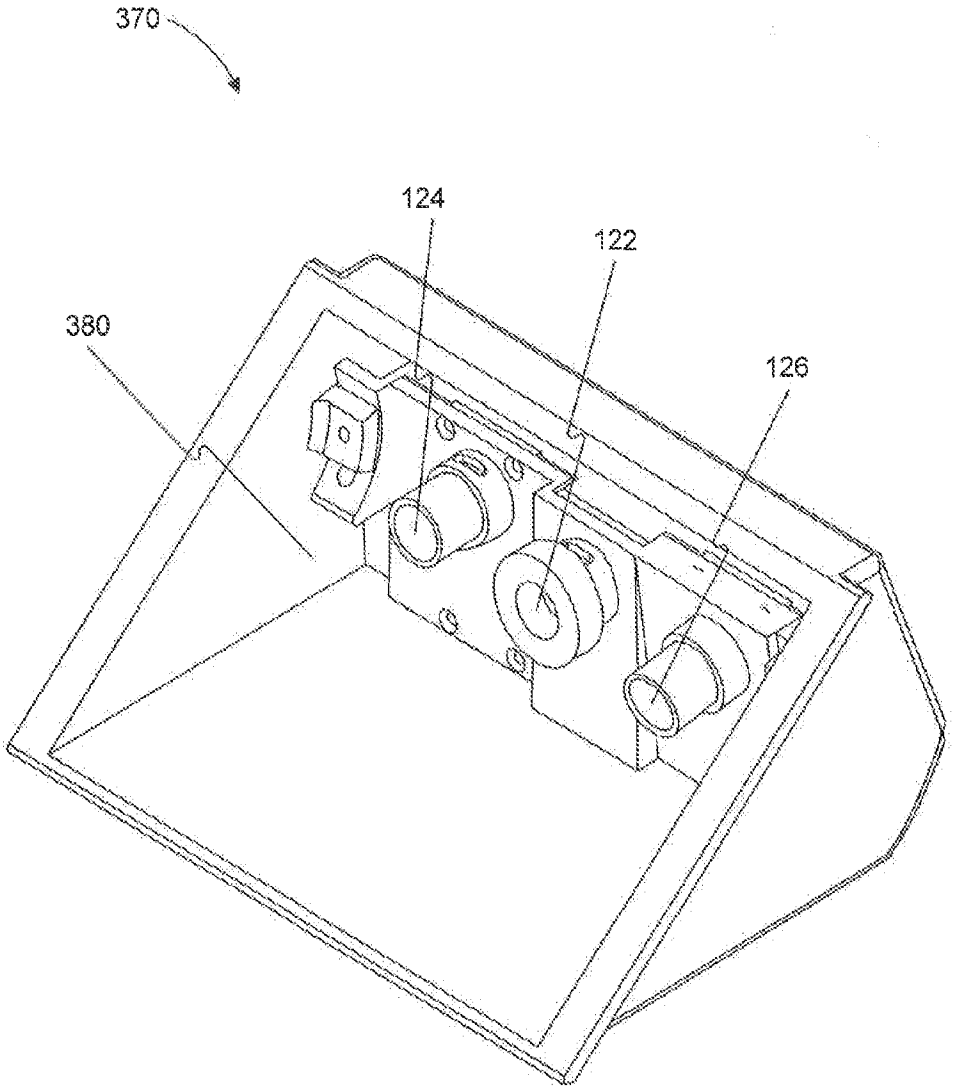


FIG. 3D

IPR2025-00943
Tesla EX1002 Page 1741

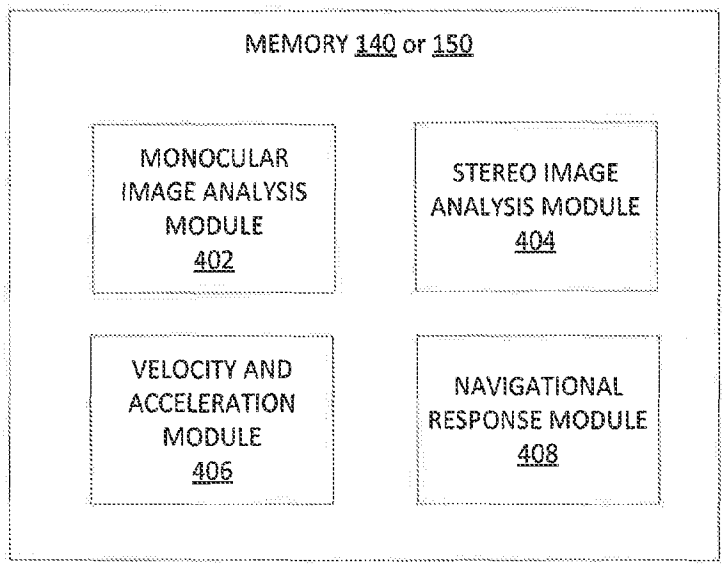


FIG. 4

IPR2025-00943
Tesla EX1002 Page 1742

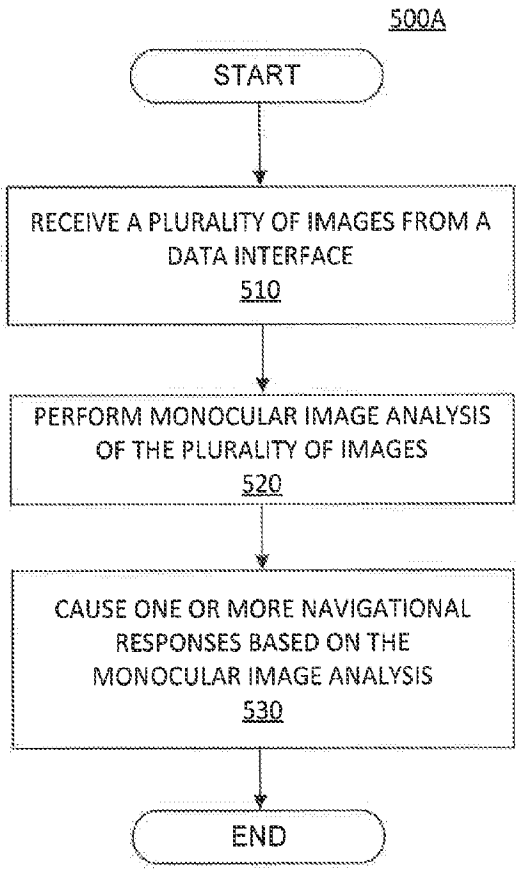


FIG. 5A

IPR2025-00943
Tesla EX1002 Page 1743

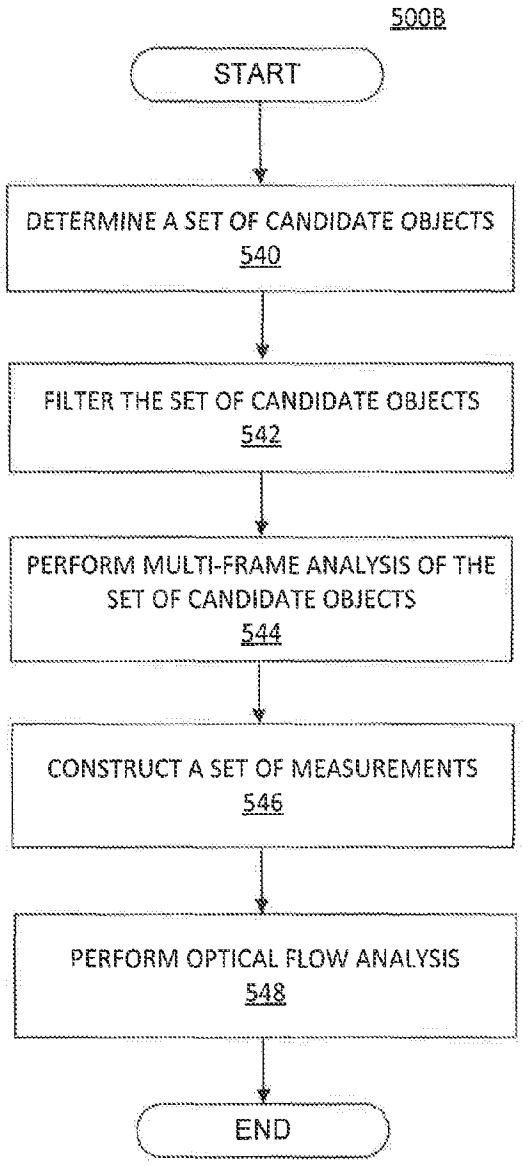


FIG. 5B

IPR2025-00943
Tesla EX1002 Page 1744

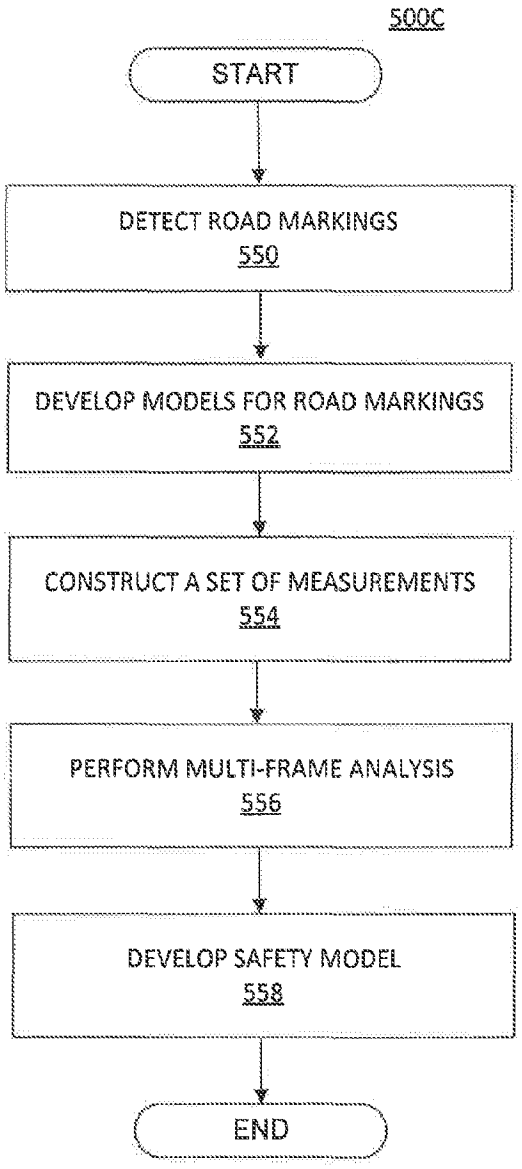


FIG. 5C

IPR2025-00943
Tesla EX1002 Page 1745

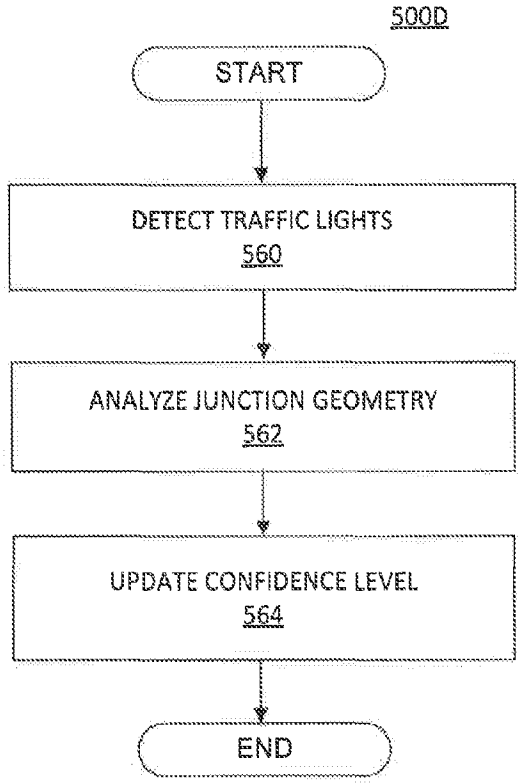


FIG. 5D

IPR2025-00943
Tesla EX1002 Page 1746

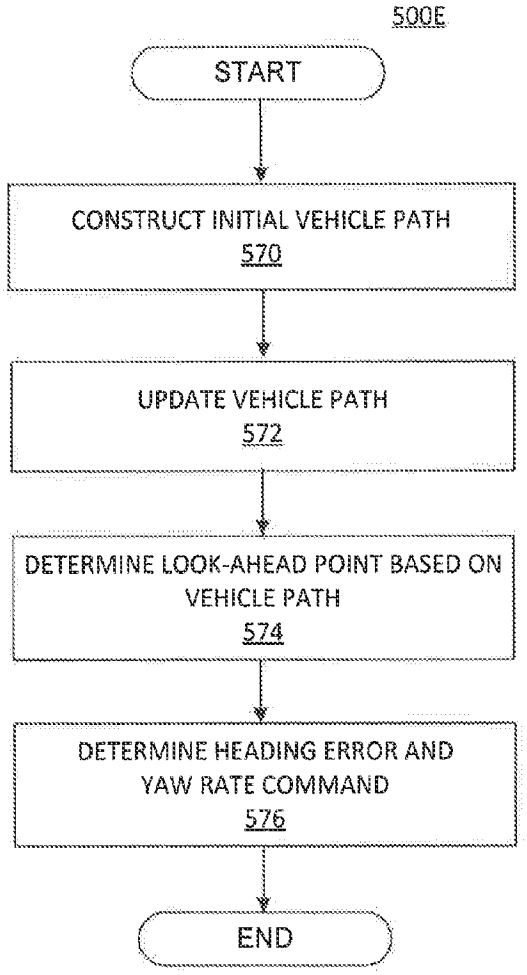


FIG. 5E

IPR2025-00943
Tesla EX1002 Page 1747

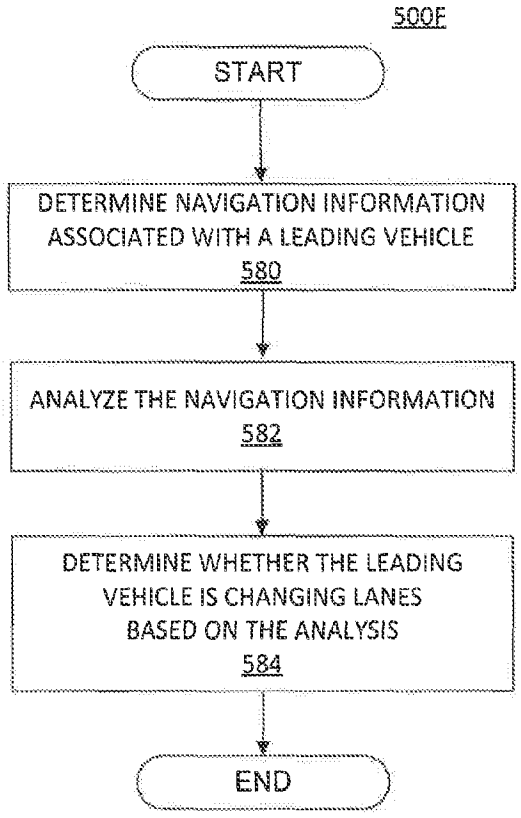


FIG. 5F

IPR2025-00943
Tesla EX1002 Page 1748

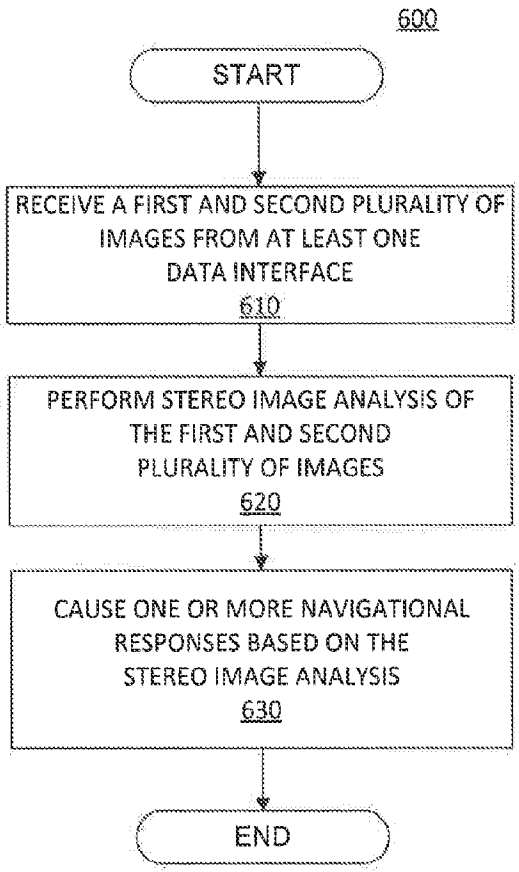


FIG. 6

IPR2025-00943
Tesla EX1002 Page 1749

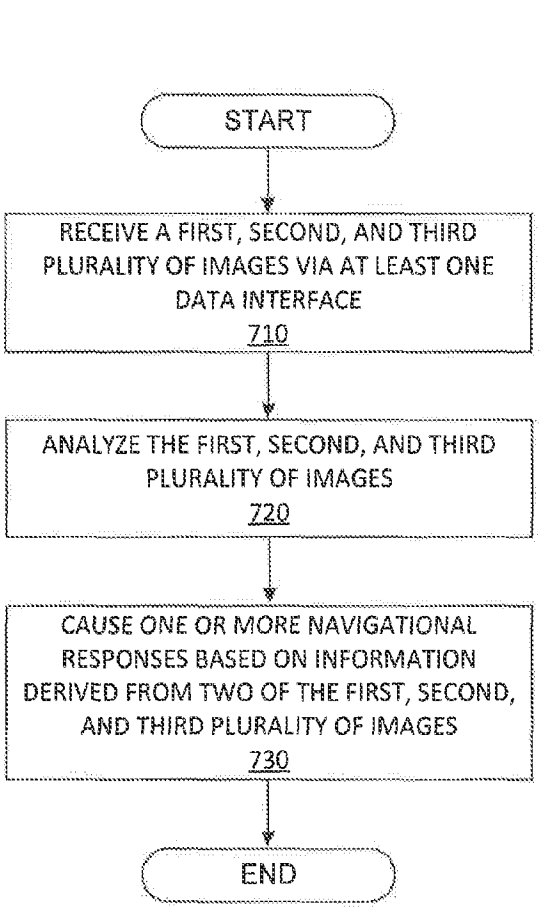


FIG. 7

IPR2025-00943
Tesla EX1002 Page 1750

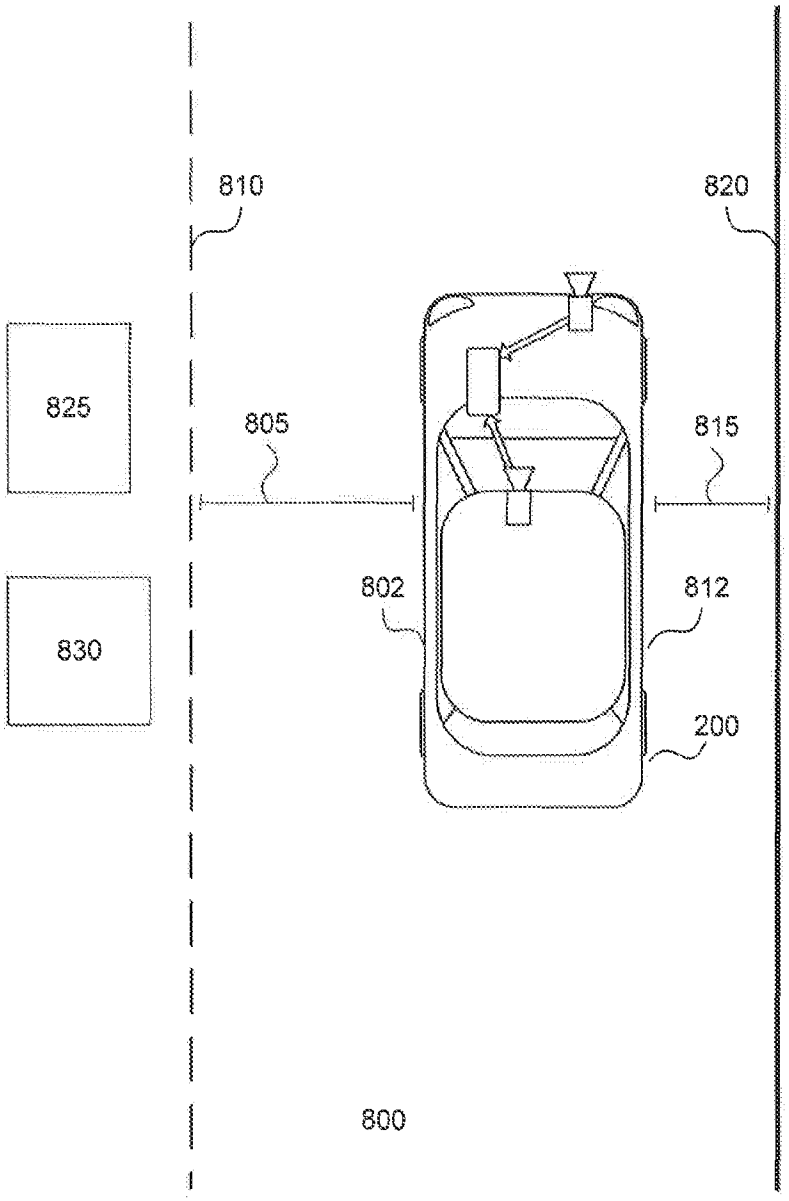


FIG. 8A

IPR2025-00943
Tesla EX1002 Page 1751

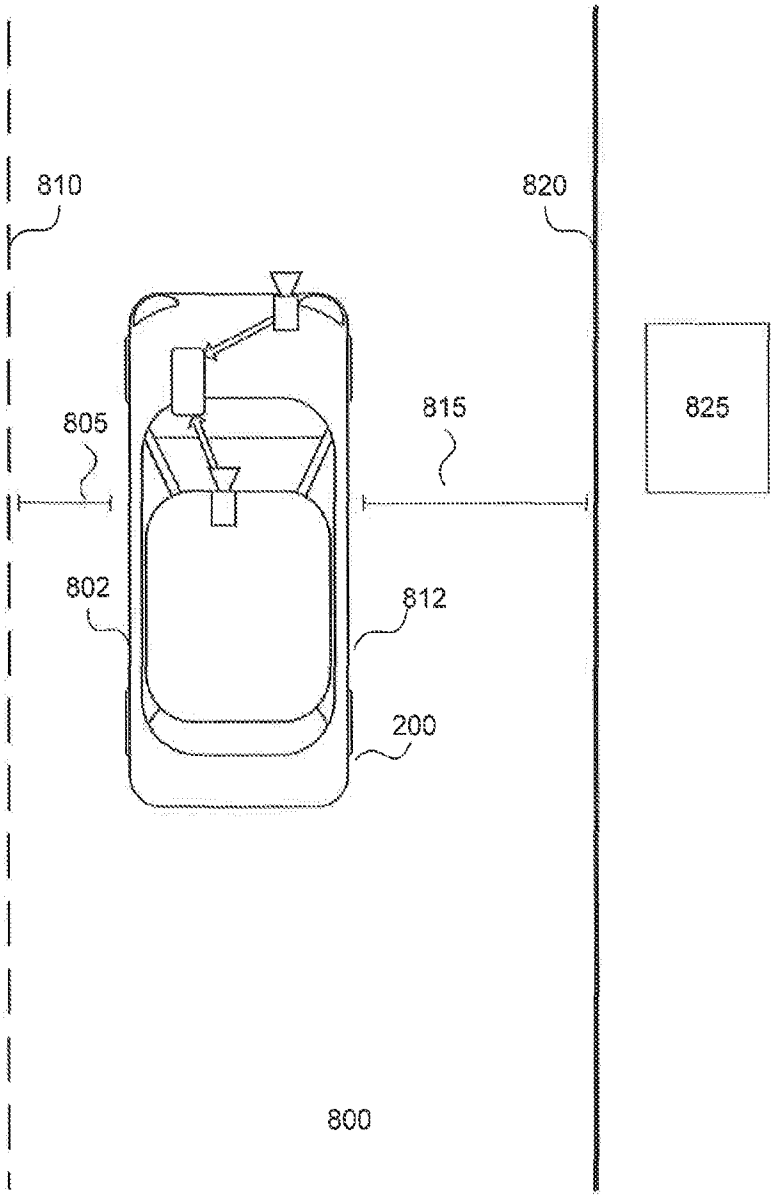


FIG. 8B

IPR2025-00943
Tesla EX1002 Page 1752

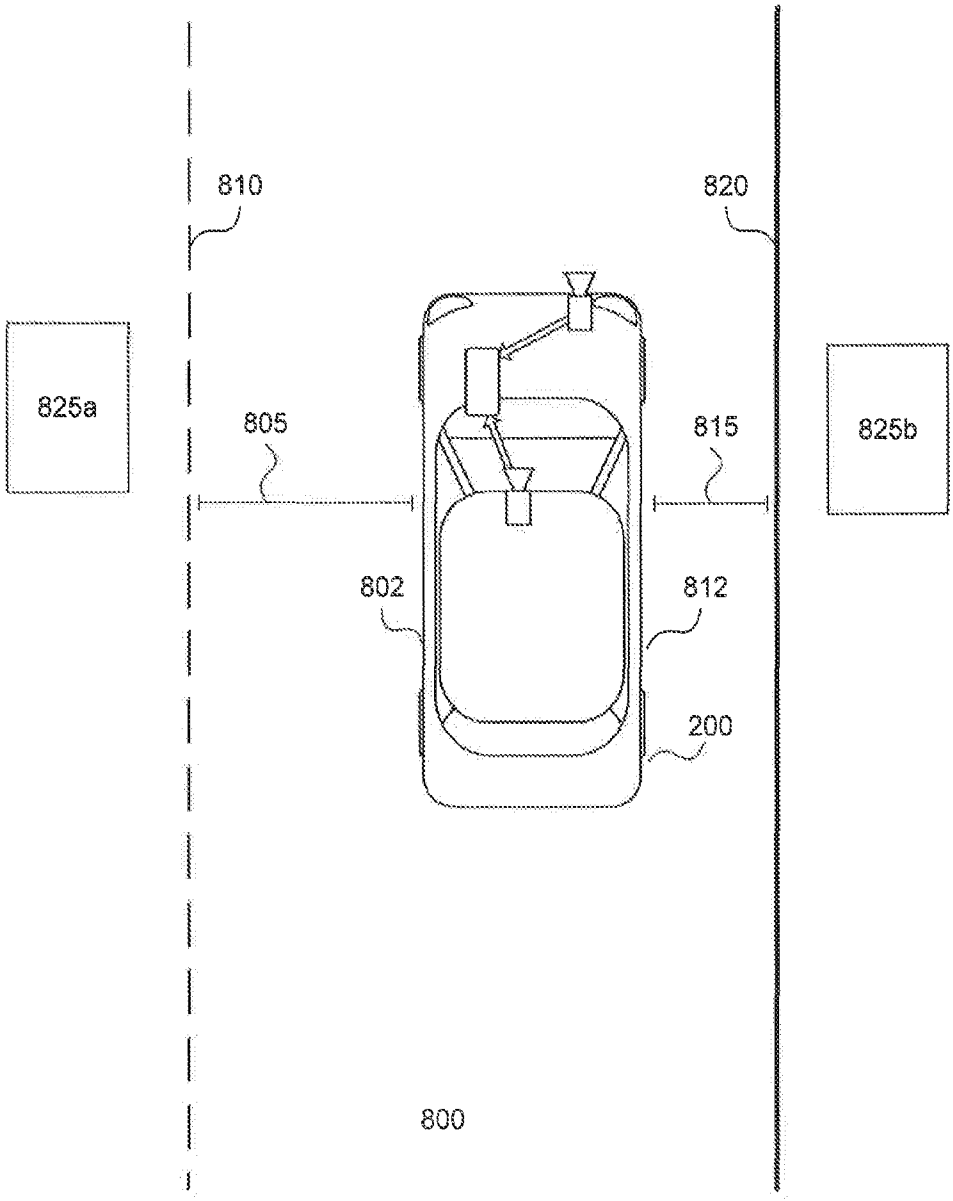


FIG. 8C

IPR2025-00943
Tesla EX1002 Page 1753

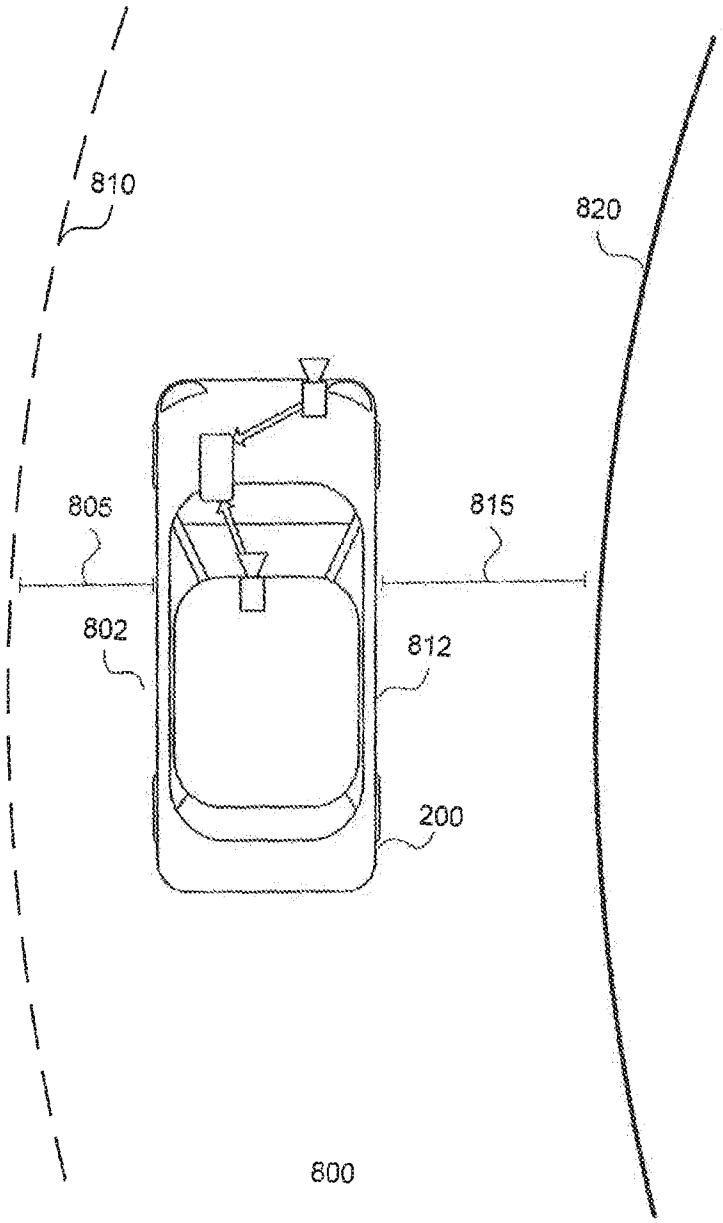


FIG. 8D

IPR2025-00943
Tesla EX1002 Page 1754

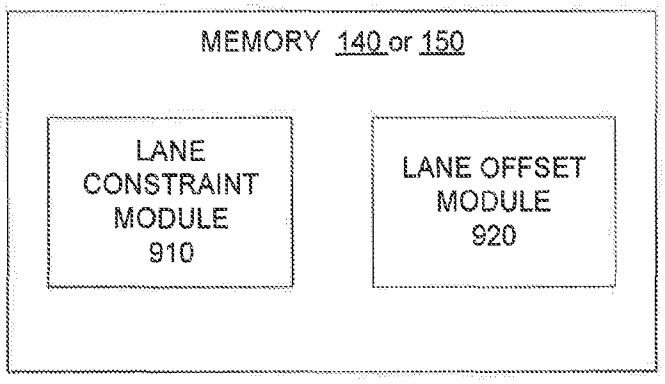


FIG. 9

IPR2025-00943
Tesla EX1002 Page 1755

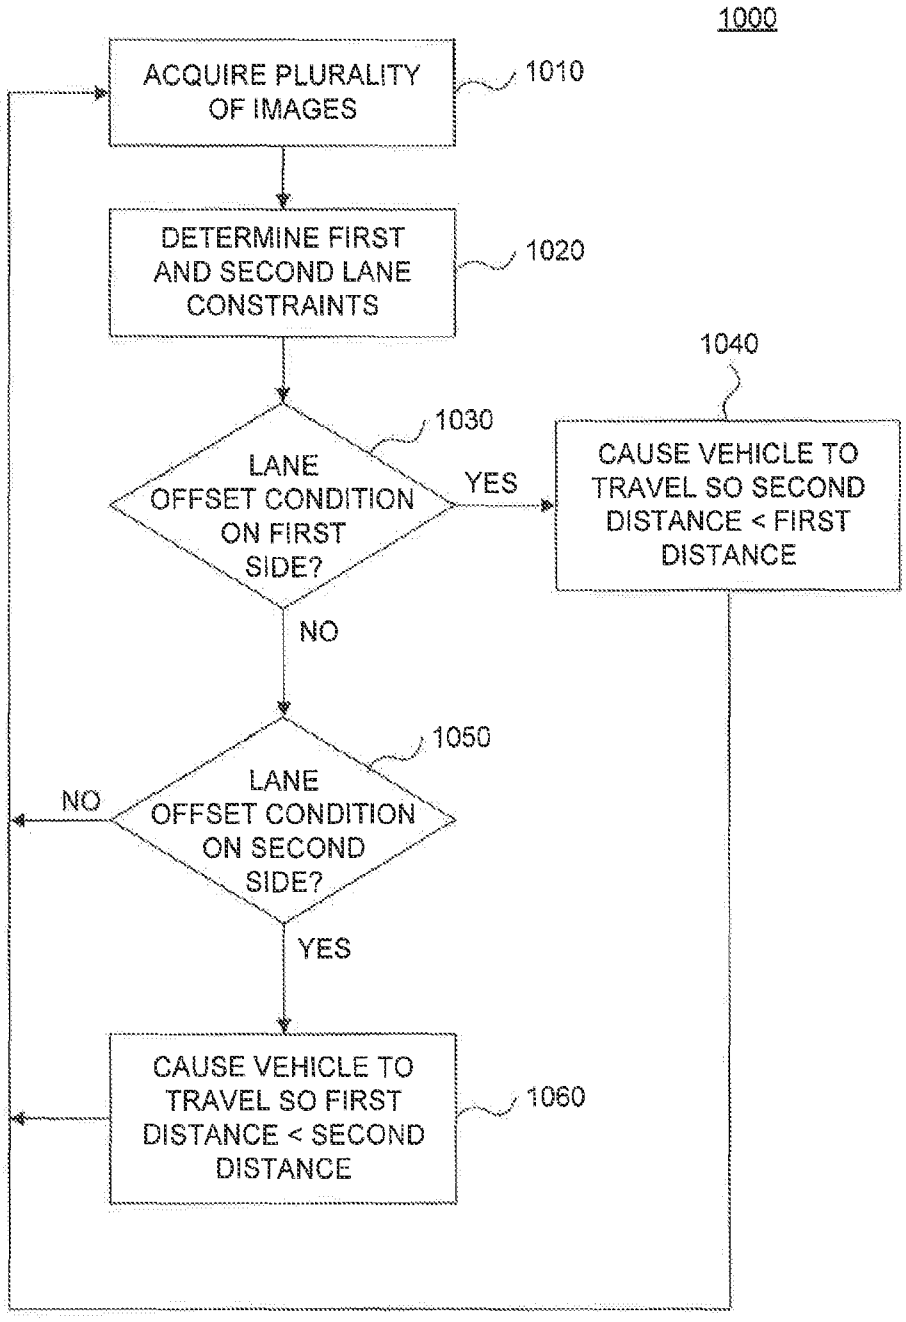


FIG. 10

IPR2025-00943
Tesla EX1002 Page 1756

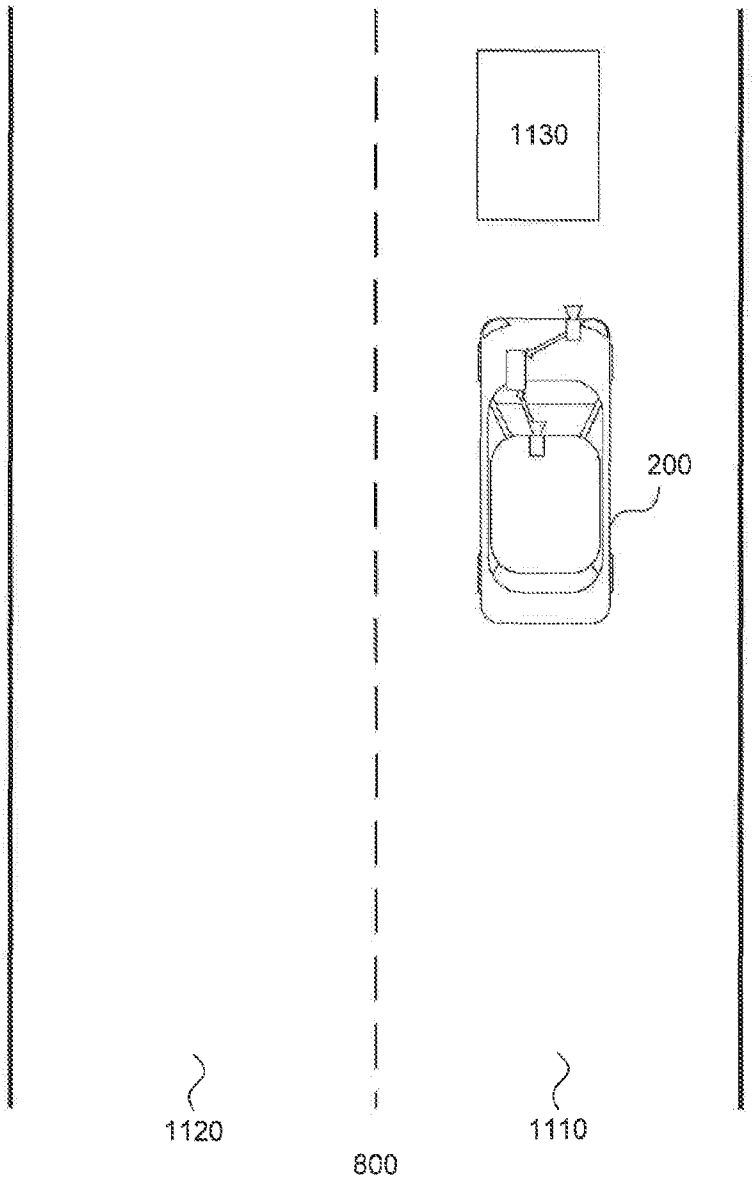


FIG. 11

IPR2025-00943
Tesla EX1002 Page 1757

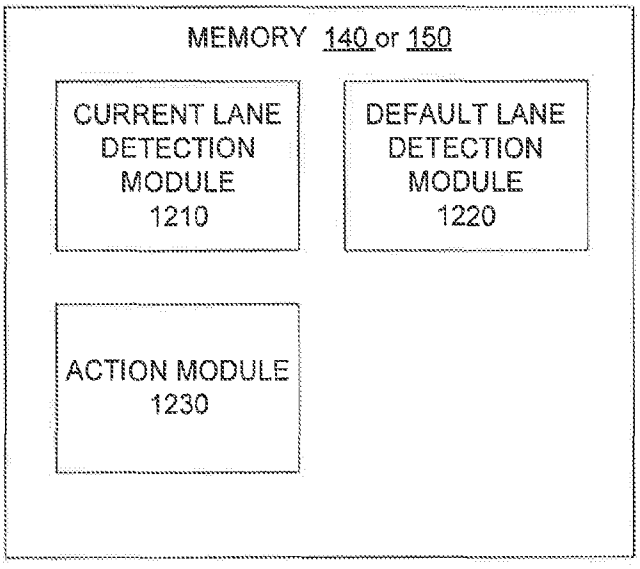


FIG. 12A

IPR2025-00943
Tesla EX1002 Page 1758

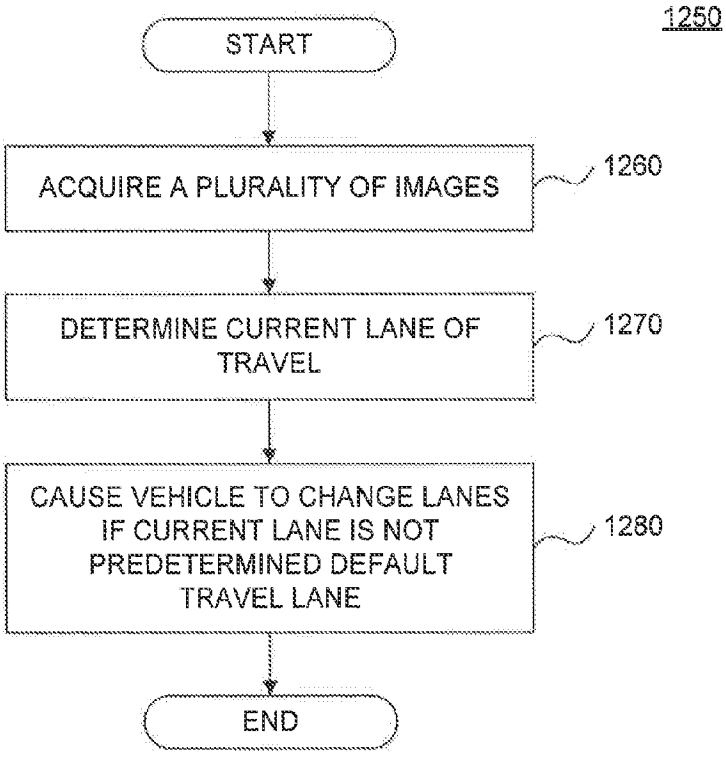


FIG. 12B

IPR2025-00943
Tesla EX1002 Page 1759

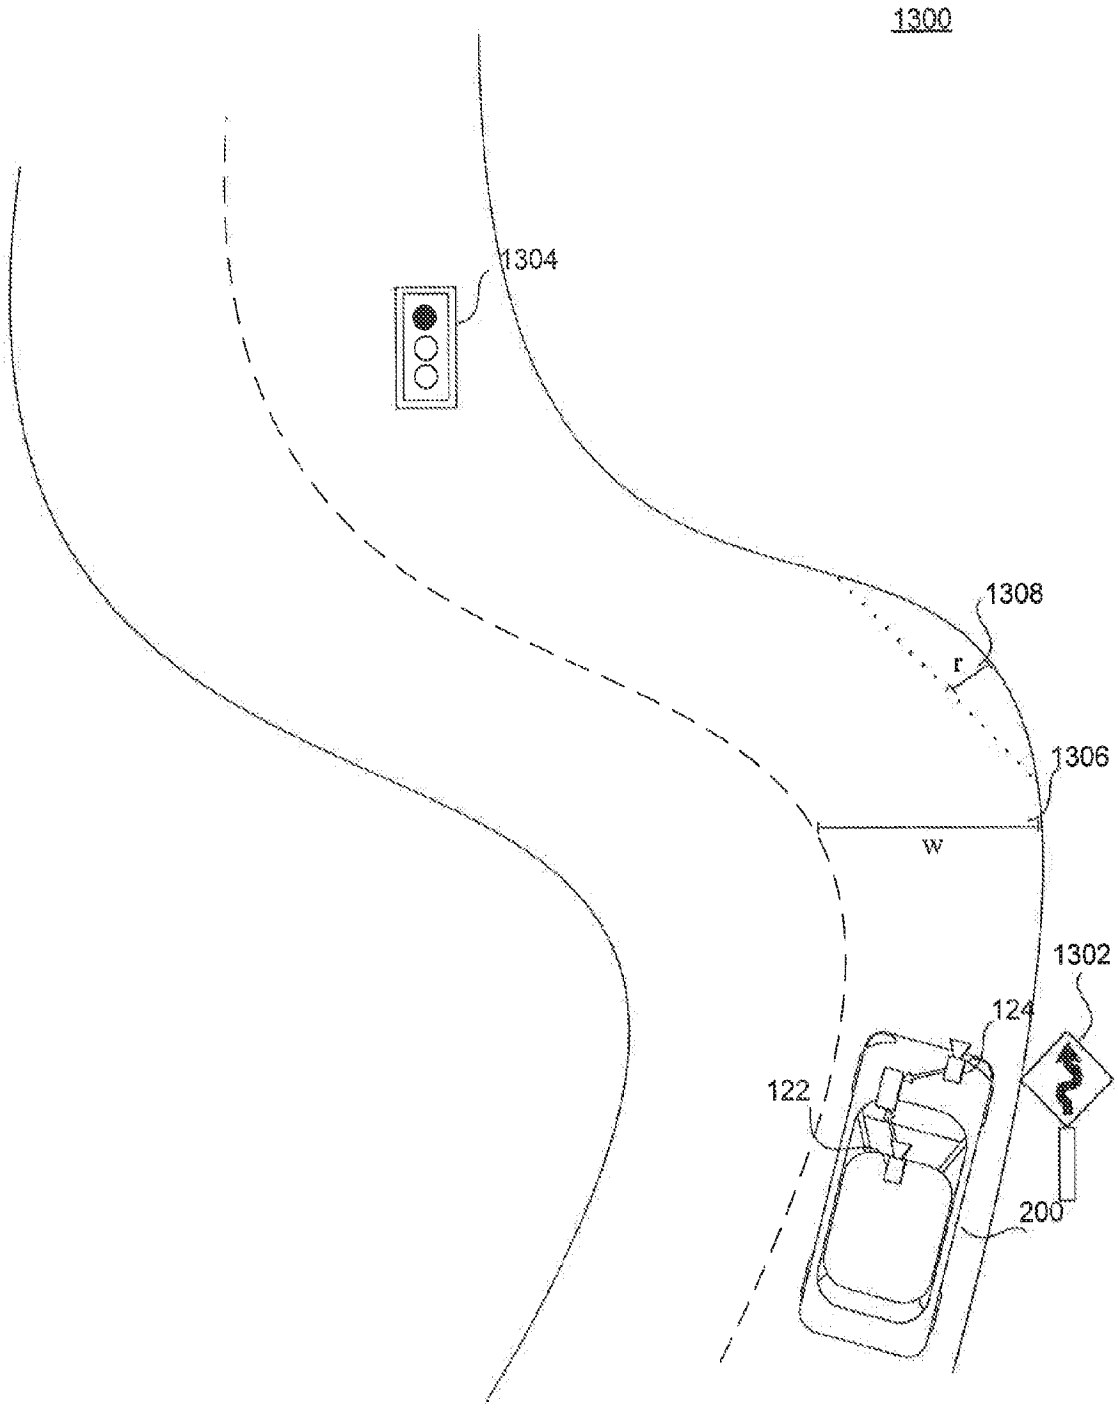


FIG. 13A

IPR2025-00943
Tesla EX1002 Page 1760

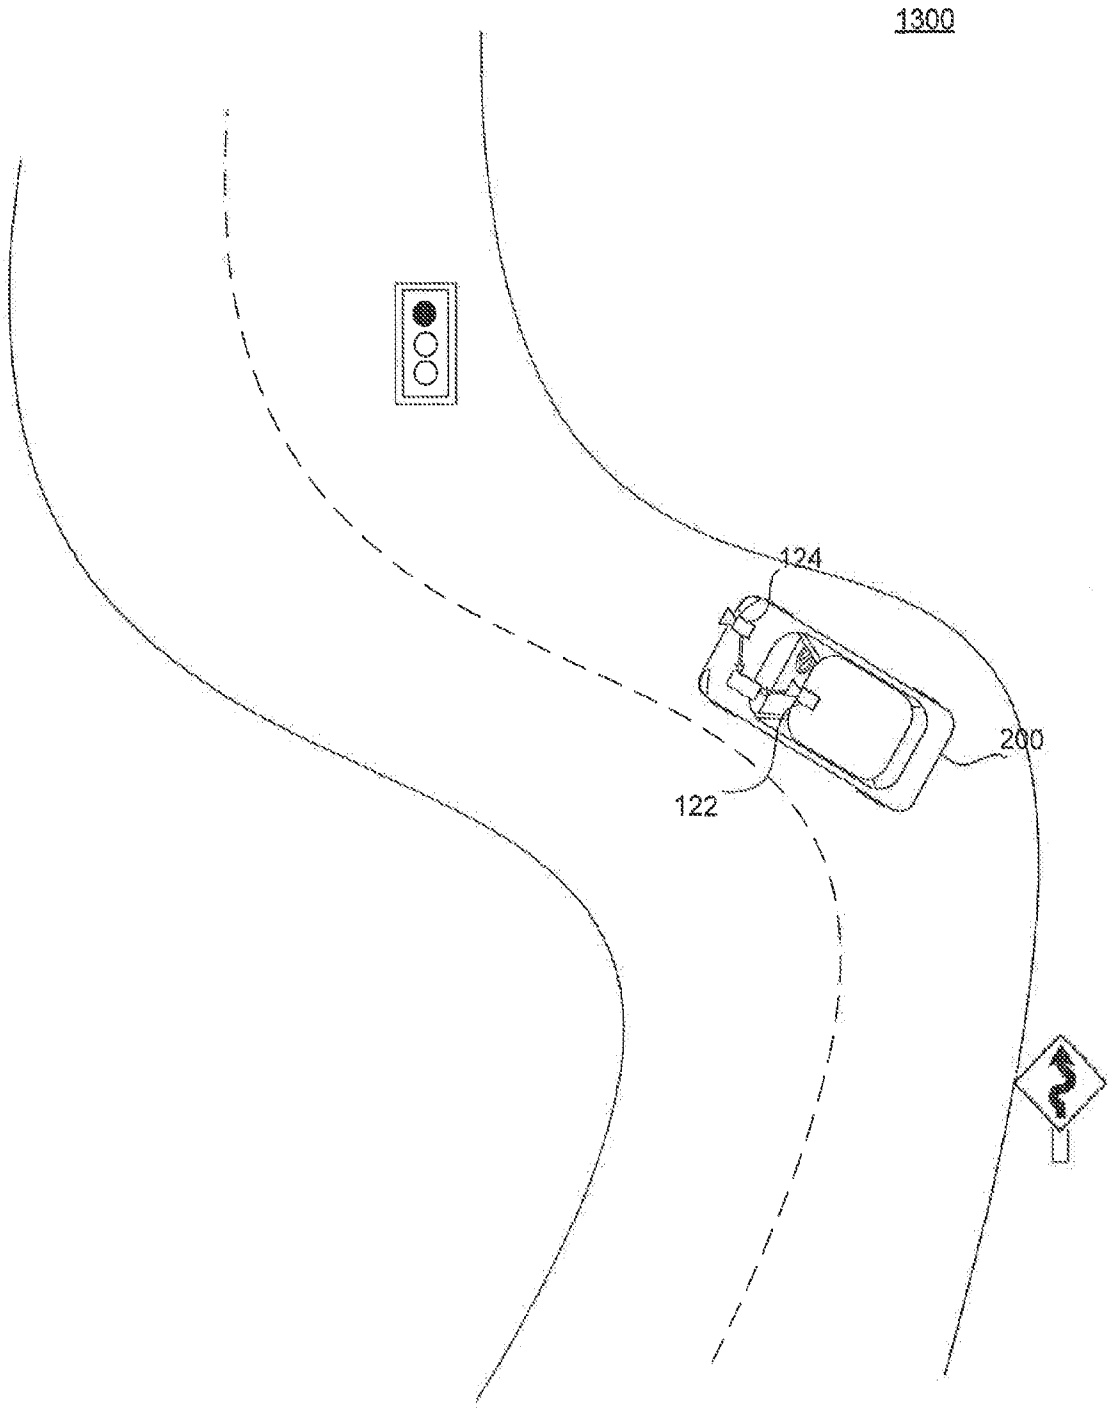


FIG. 13B

IPR2025-00943
Tesla EX1002 Page 1761

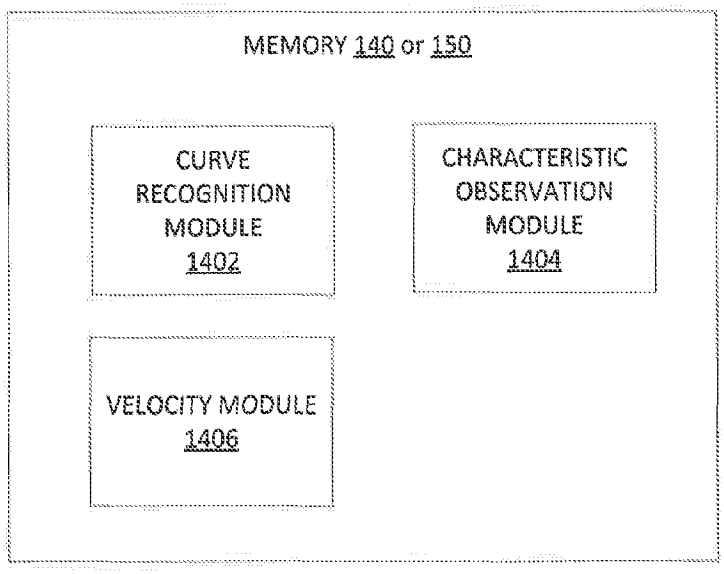


FIG. 14

IPR2025-00943
Tesla EX1002 Page 1762

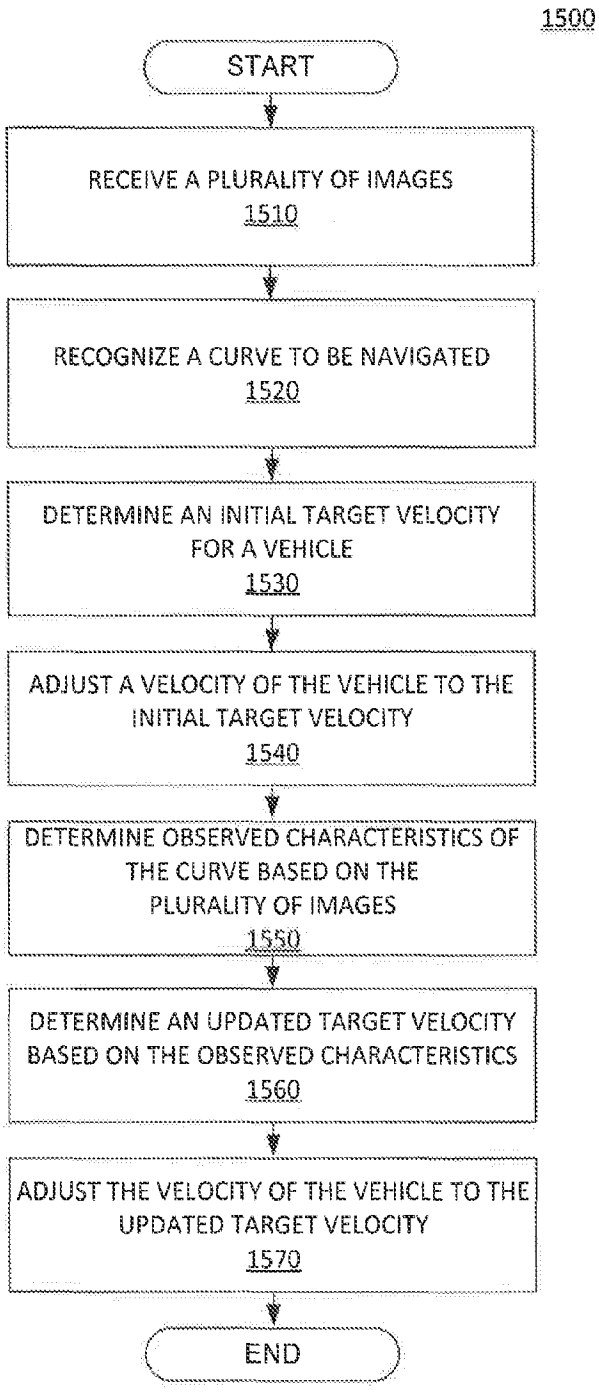


FIG. 15

IPR2025-00943
Tesla EX1002 Page 1763

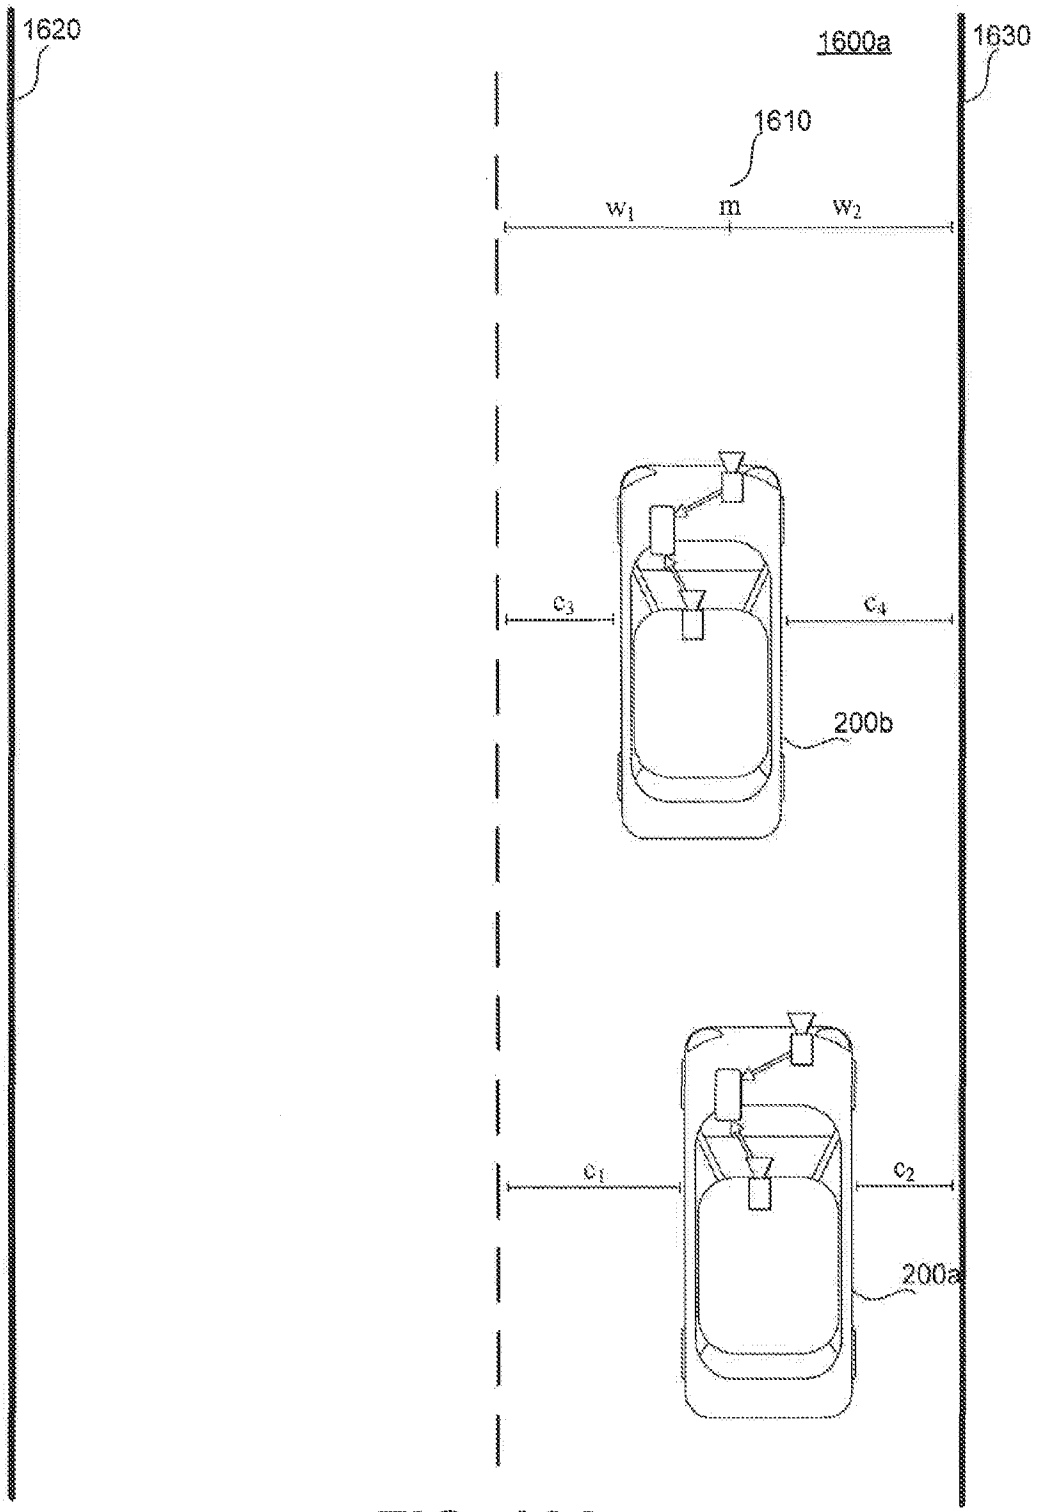


FIG. 16A

IPR2025-00943
Tesla EX1002 Page 1764

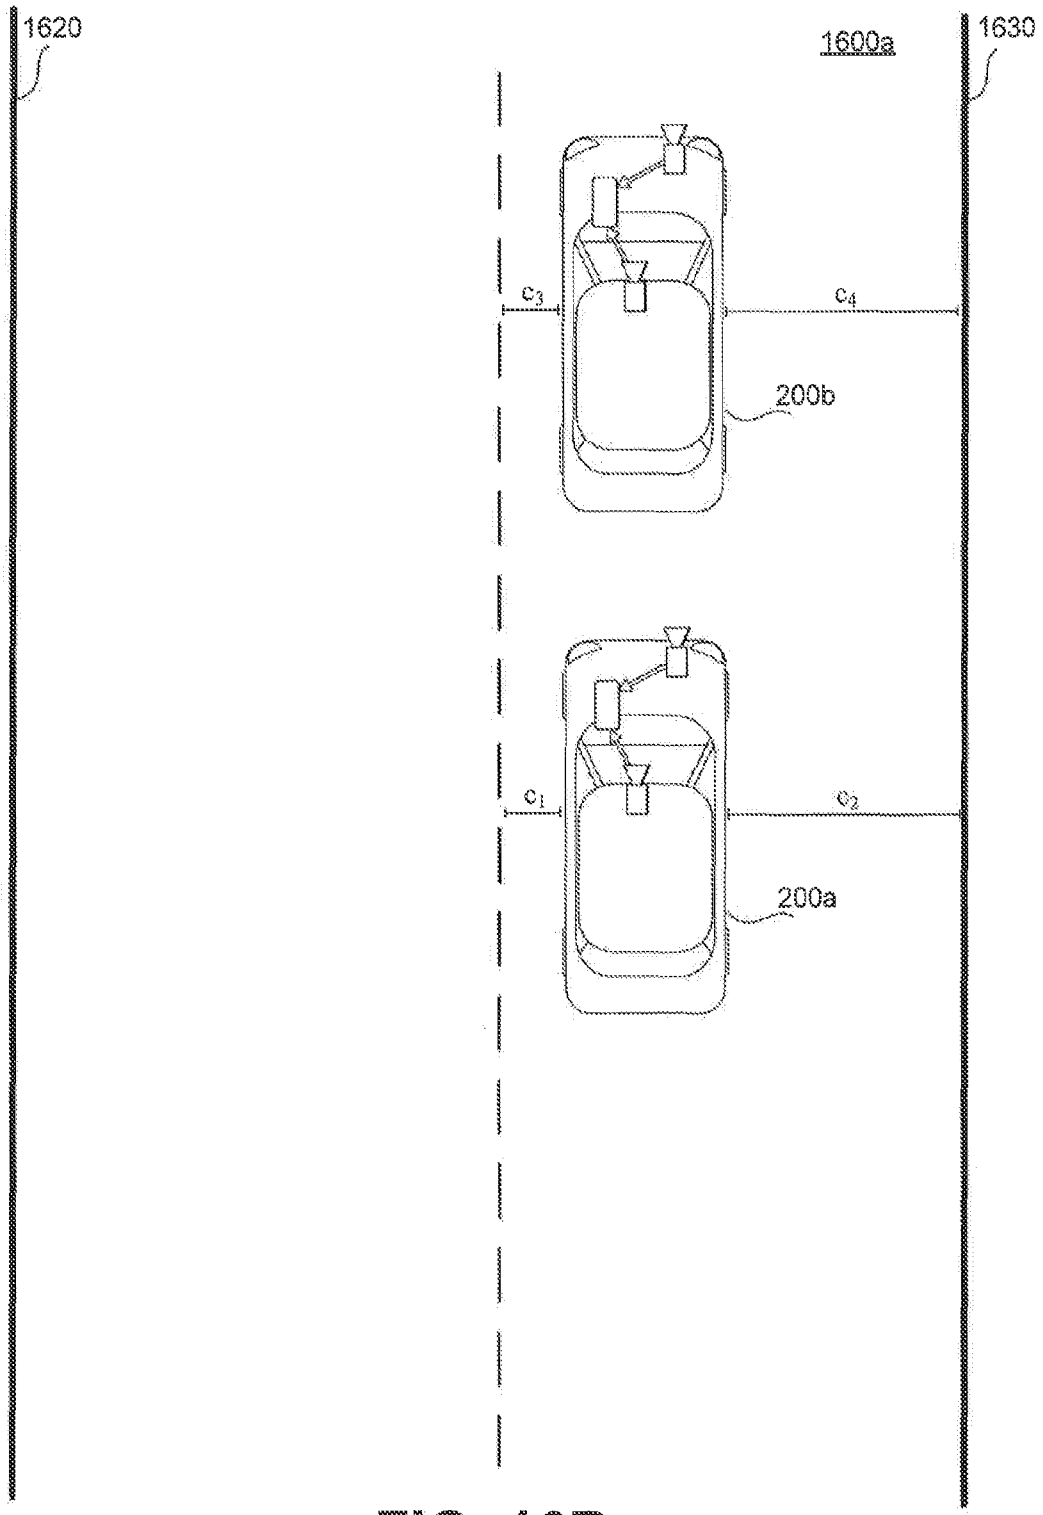


FIG. 16B

IPR2025-00943
Tesla EX1002 Page 1765

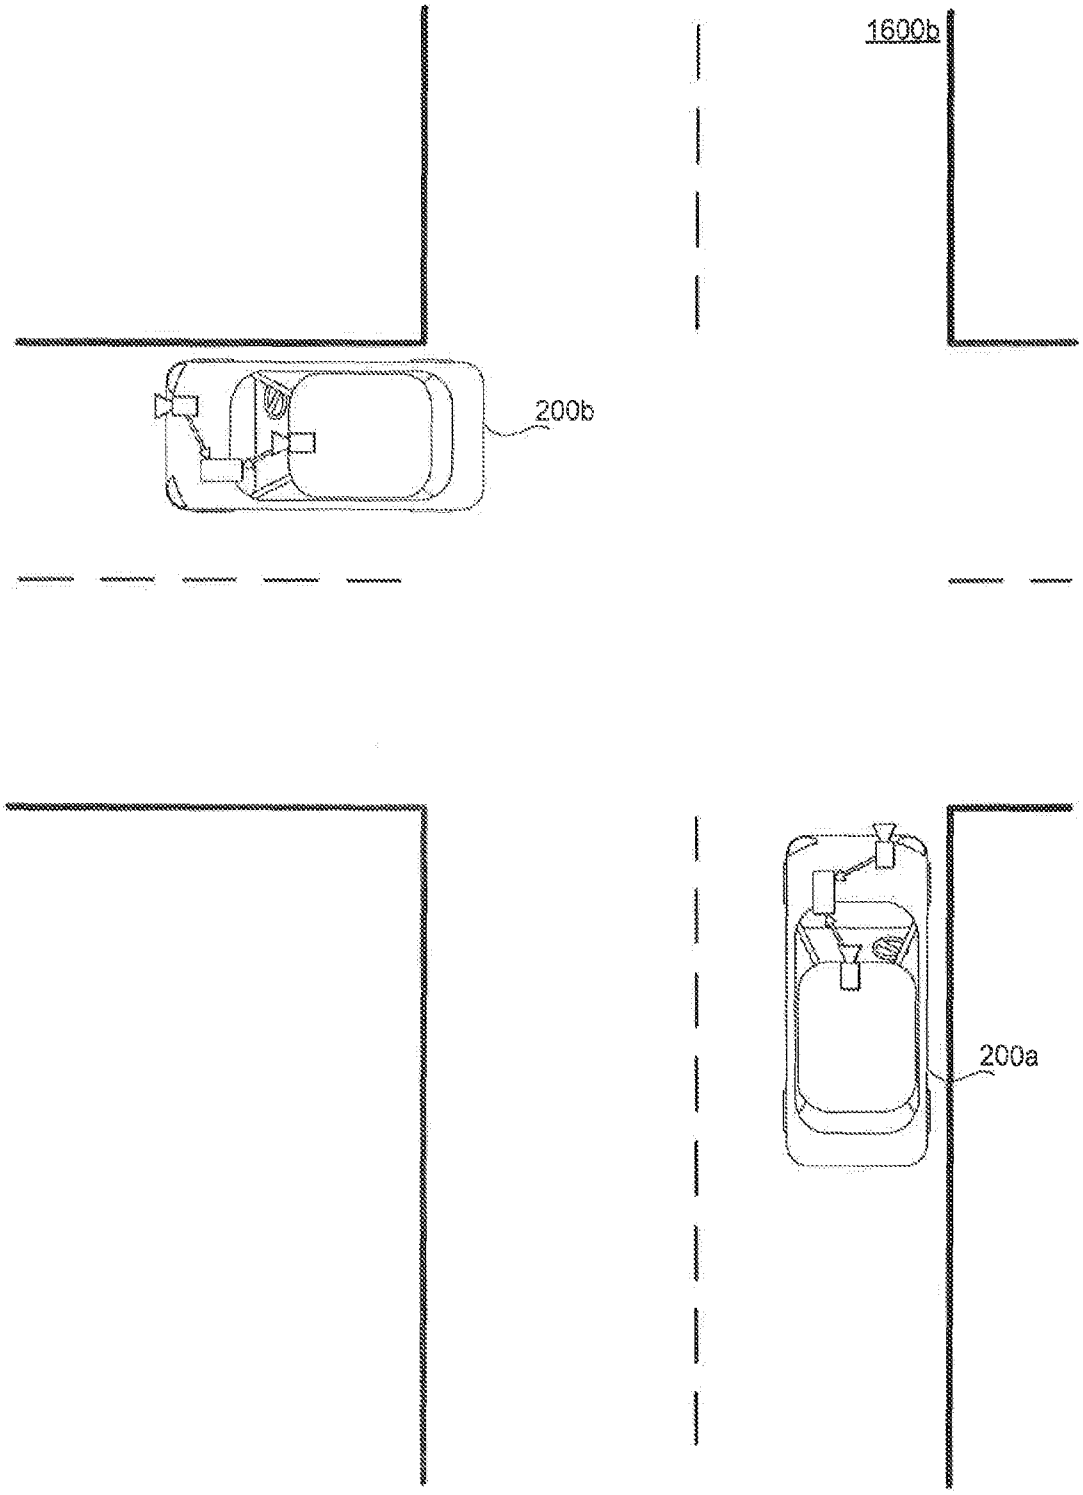


FIG. 16C

IPR2025-00943
Tesla EX1002 Page 1766

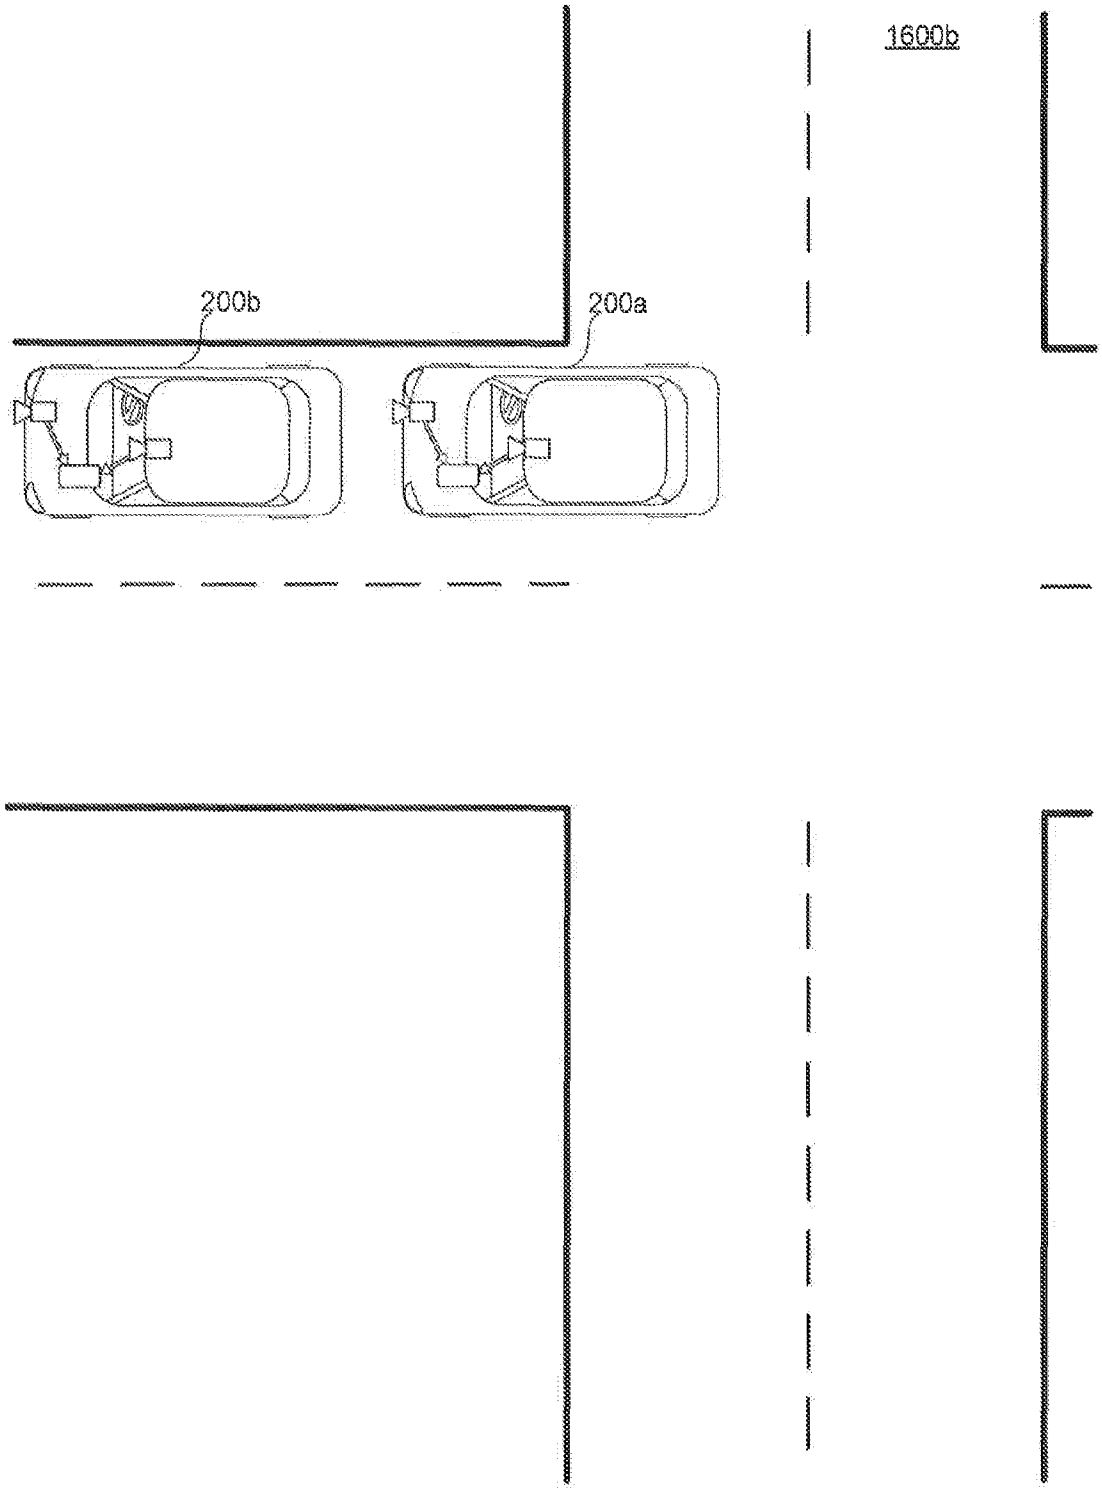


FIG. 16D

IPR2025-00943
Tesla EX1002 Page 1767

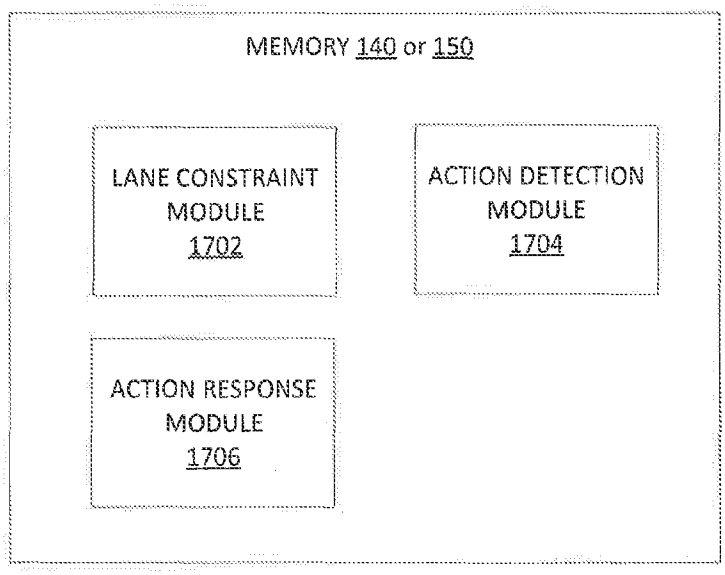


FIG. 17

IPR2025-00943
Tesla EX1002 Page 1768

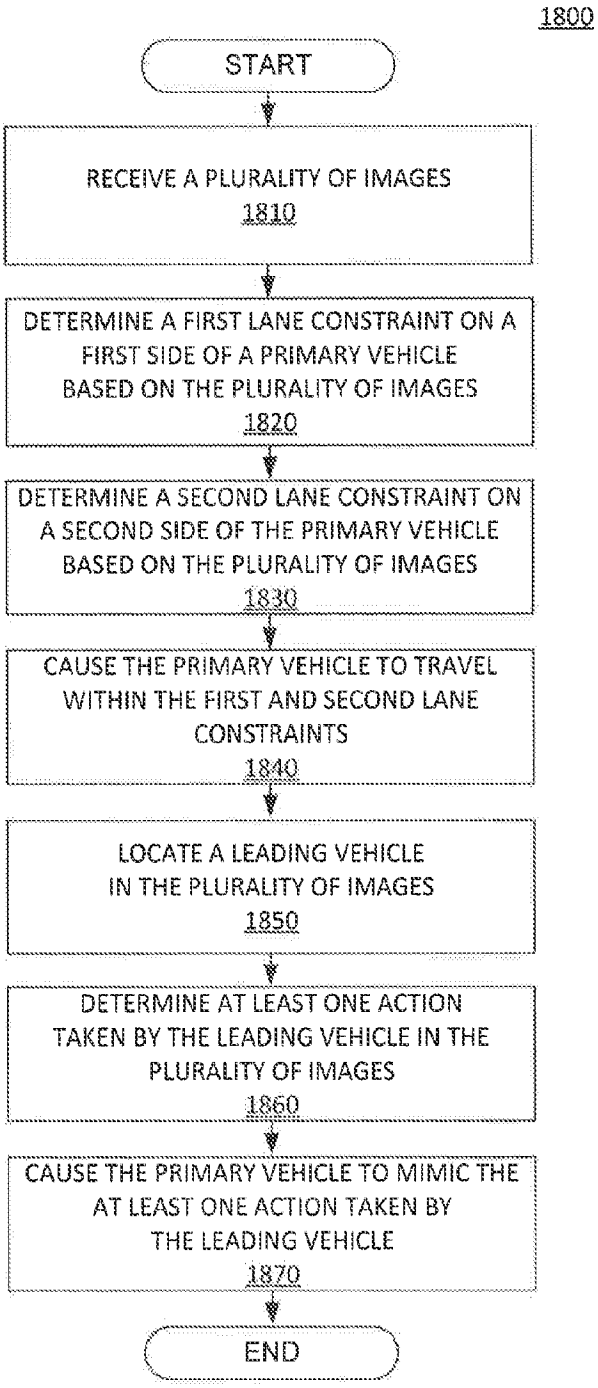


FIG. 18

IPR2025-00943
Tesla EX1002 Page 1769

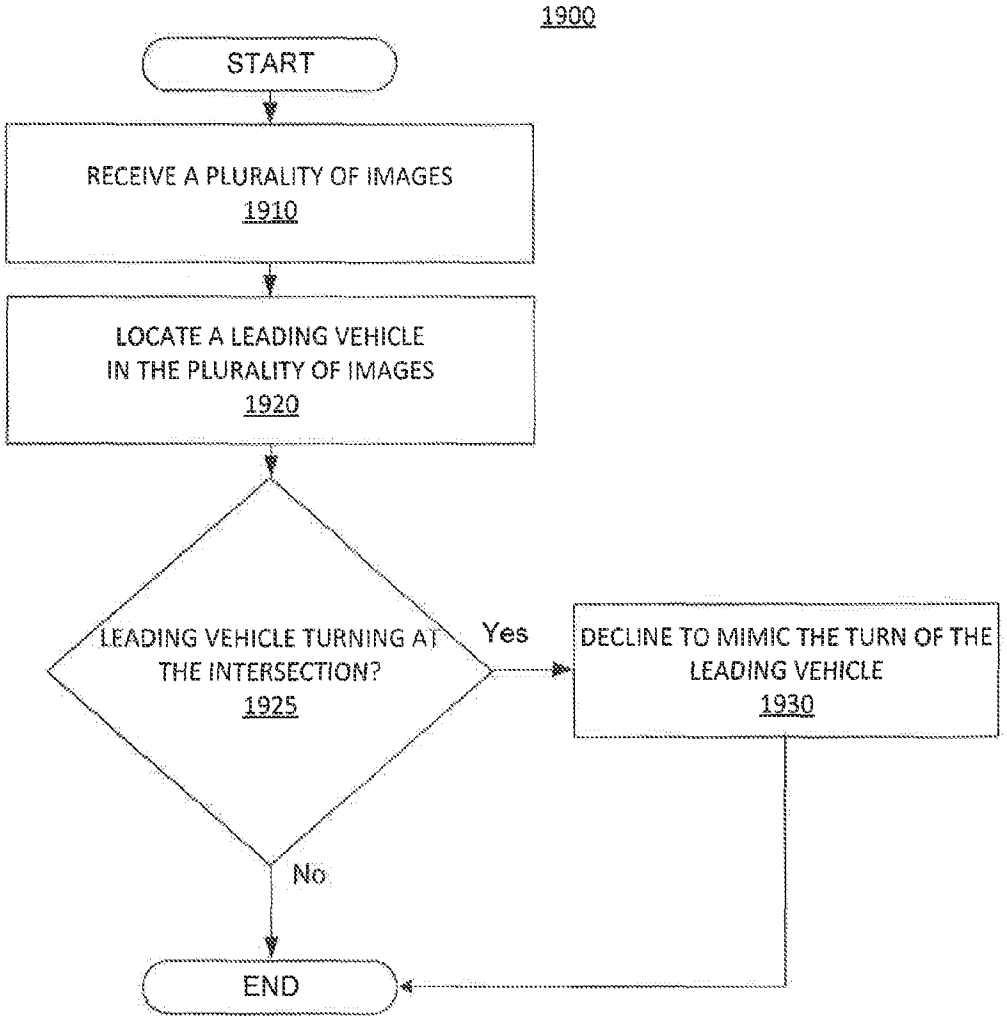


FIG. 19

IPR2025-00943
Tesla EX1002 Page 1770

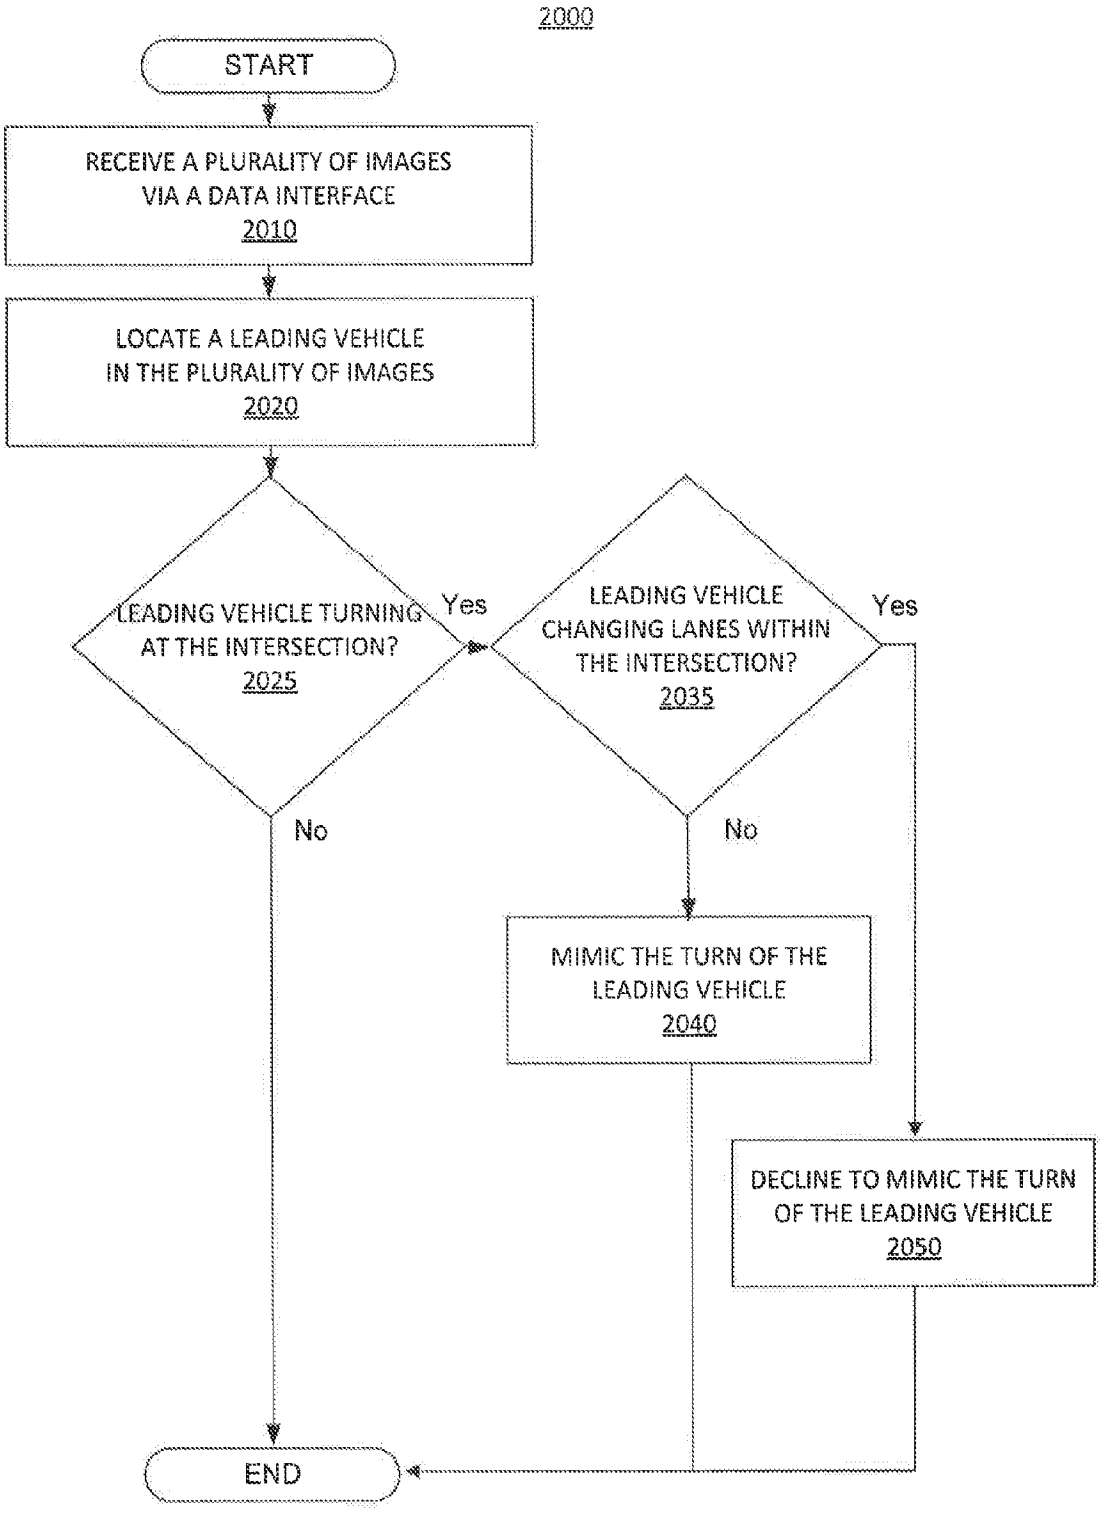


FIG. 20

IPR2025-00943
Tesla EX1002 Page 1771