

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSARA, INC.

Petitioner

v.

MOTIVE TECHNOLOGIES, INC.,

Patent Owner.

Case IPR2025-00574

U.S. Patent No. 11,875,580

DECLARATION OF SCOTT ANDREWS

I, Scott Andrews, do hereby declare and state as follows:

I. BACKGROUND AND QUALIFICATIONS

1. I am currently a consultant for Cogenia Partners, LLC, focusing on systems engineering, business development and technical strategy supporting automotive and information technology. I have been in this position since 2001.

2. I have acted in various roles as a technical expert and consultant in many areas involving automotive systems. For example, I was a technical consultant for sensor supplemented message validation system for vehicle-to-vehicle communication-based collision warning/avoidance systems in a project sponsored by the U.S. Dept. of Transportation's National Highway Traffic Safety Administration (NHTSA). I served as the technical lead on a project to develop requirements for connected vehicle safety systems in preparation for NHTSA regulations governing such systems. I also serve as a technical consultant on multiple projects sponsored by the Federal Highway Administration (FHWA) related to connected vehicle technology research.

3. I have over 30 years of professional experience in the field automotive technologies and systems, including vehicle sensor systems and vehicle safety and control systems. Further, I have authored numerous published technical papers and am a named inventor on 11 U.S. and foreign patents.

4. I received a Bachelor of Science degree in Electrical Engineering from University of California, Irvine in 1977 and a Master of Science degree in Electronic Engineering from Stanford University in 1982.

5. From 1977 to 1979, I worked at Ford Aerospace where I designed, tested and delivered microwave radar receiver system

6. From 1979 to 1983, I worked at Teledyne Microwave, where I developed high reliability microwave components and developed CAD tools. From 1983 to 1996, I worked at TRW, Inc., having held various positions.

7. From 1983 to 1985, I was a Member of the technical staff and a Department Manager in the Space Electronics sector. Between 1985 and 1990 I was a project manager working on various communications systems projects including the US DoD Advanced Research Projects Administration (ARPA) MIMIC Program. Between 1990 and 1993 I was the Manager of MMIC (monolithic-microwave-integrated-circuit) Products Organization. In this role, I developed business strategy and managed customer and R&D programs. During this time, I also developed the first single chip 94 GHz Radar, used for automotive cruise control and anti-collision systems. In 1993 I transferred to the TRW Automotive Electronics Group, and managed about 30 engineers in the Systems Engineering and Advanced Product Development organization. In this role, I managed advanced development programs such as automotive radar, adaptive

cruise control, occupant sensing, automatic crash notification systems, in-vehicle information systems, and other emerging transportation products.

8. I was employed as a Project General Manager in the Electronics Division of Toyota Motor Corporation. I worked at Toyota headquarters in Toyota City, Japan from April 1996 to around April 2000. Between July 1999 and April 2000, I transitioned from working in Japan to working in a Toyota office in San Jose, CA. In this position, I was responsible for leading the development of vehicle telematics systems, infotainment systems, including on-board and off-board navigation systems, traffic information systems, vehicle communications systems, safety applications, and automated vehicle control systems.

9. In 1998, I founded the Automotive Multimedia Interface Collaboration, a consortium of car makers developing standards for in-vehicle computing and interfaces between consumer multimedia systems and consumer electronics devices. This work resulted in a variety of standards for vehicle interfaces, user interfaces and vehicle software management that were eventually transferred to other standards organizations such as ISO and the OSGi Alliance.

10. In the various positions mentioned above, I was responsible for research and development projects relating to numerous vehicle information systems, user interface systems, sensory systems, control systems and safety systems, and also had the opportunity to collaborate with numerous researchers

and suppliers to the auto industry. I therefore believe that I have a detailed understanding of the state of the art during the relevant period, as well as a sound basis for opining how persons of skill in the art at that time would understand the technical issues in this case. In 2000, I founded Cogenia, Inc. to develop enterprise class data management software systems. I served as the company's Chief Executive Officer until 2001, when I created Cogenia Partners, my current consulting firm.

11. A copy of my *curriculum vitae* is attached hereto, and it includes a listing of my prior experience in litigation matters as an expert.

II. ASSIGNMENT AND MATERIALS CONSIDERED

11. I submit this declaration in support of Motive Technologies Inc.'s Patent Owner Preliminary Response to Samsara Inc.'s Petition for *inter partes* review of U.S. Patent No. 11,875,580 ("the '580 patent").

12. I am not an employee of Motive Technologies or of any affiliate or subsidiary thereof.

13. My consulting firm, Cogenia Partners, LLC, is being compensated for my time at a rate of \$600 per hour.

14. My compensation is in no way dependent upon the substance of the opinions I offer below, or upon the outcome of the petition for *inter partes* review.

15. I have been asked to provide certain opinions relating to the

patentability of the '580 patent and the viability of the grounds presented in the Petition. Specifically, I have been asked to provide my opinion regarding (i) the level of ordinary skill in the art to which the '580 patent pertains, (ii) whether any grounds in the Petition show any basis to find claims 1-7 of the '580 patent unpatentable as being obvious in view of the prior art asserted in the Petition.

16. Given the preliminary nature of the Patent Owner response, the opinions expressed in this declaration are not exhaustive of my opinions on the patentability of claims 1-7 of the '580 patent. Therefore, the fact that I do not address a particular point should not be understood to indicate any opinion on my part that any arguments or grounds advanced in the Petition are meritorious.

17. In forming my opinions, I have reviewed the Petition, including its exhibits, the '580 patent disclosure (including materials incorporated by reference) and its prosecution history, as well as the following patents and publications:

- a) U.S. Patent Application Publication No. 2019/0034740 to Kwant *et al.* ("Kwant1") (Ex. 1057)
- b) U.S. Patent Application Publication No. 2019/0102674 to Kwant *et al.* ("Kwant2") (Ex. 1058)
- c) U.S. Publication No. 2020/0410704 to Choe *et al.* ("Choe") (Ex. 1005)
- d) Khan *et al.*, "Lane Detection Using Lane Boundary Marker

Network with Road Geometry Constraints,” (“Khan”) (Ex. 1007)

e) Ali et al., “Real-Time Vehicle Distance Estimation using Single View Geometry,” (“Ali”) (Ex. 1008)

f) PCT Application Publication No. PCT/EP2019/056356 to Westmacot *et al.* (“Westmacot”) (Ex. 1010)

g) Scott Workman and Menghua Zhai and Nathan Jacobs, “Horizon Lines in the Wild,” (“Workman”) (Ex. 1011)

III. OVERVIEW OF THE ’580 PATENT AND ITS PROSECUTION

18. The inventors of the ’580 Patent identified an issue with the use of dash-cams in vehicles. It is common to use automotive sensors for detection of objects on the road, but typically “[m]ultiple cameras or sensors” are needed. Ex. 1001 at 1:10-11. The “use of multiple cameras” or proximity sensors “such as Radar and Lidar” was too expensive. *Id.* at 1:11-14. Monocular cameras provide a more cost-effective approach, but their use “require[d] the location of the camera to be predefined or preset.” *Id.* at 14-16. The inventors resolve this problem with a method to re-initialize the camera upon relocation. *Id.* at 1:28-29.

19. The ’580 patent describes embodiments for “using a monocular camera (e.g., an image sensor) that may be retrofitted and adjustable within the vehicle such as, the vehicle’s dashboard.” *Id.* at 1:21-23. “Upon driving, the camera initializes itself to determine its height with respect to the ground as well as

a road plane normal. Thus, the camera may be installed in various positions and later re-initialized or recalibrated as needed.” *Id.* at 1:24-27.

20. The method may, for example, be used to provide “a network-based initialization system.” *See e.g., id.* at 1:36-52. The specification describes that:

[A] server-side system receives a video comprising a set of image frames and identifies one or more lines in the video using a predictive model (e.g., a convolutional neural network). In one embodiment, the lines can include a horizon line. After identifying the horizon line, the system computes some camera parameters (e.g., height, viewing angle, road plane normal, etc.) based on the horizon line. The system then overlays the computed horizon line on the video and transmits it to an annotator device for manual review. In response, the annotator device can transmit a confirmation indicating the horizon line was accurate. In some embodiments, the annotator device can reject the horizon line and manually add the horizon line. Once the horizon line is confirmed, the system returns the camera parameters to the calling party (e.g., a camera onboard a vehicle).

Id. Such a network-based camera initialization service provides a robust approach to initializing dash-cams in a reliable way using “a remote computing device executing the camera initialization service 232,” for example, “an application running on a remote server.” *See e.g., id.* at 5:13-20.

21. The specification explains that the disclosed “[s]ystems can use the illustrated neural network in some embodiments instead of a geometric process to determine lane and horizon lines.” *Id.* at 7:14-16. “Specifically, given an input image 302 (e.g., a video frame), a backbone network 304 can convolve the input image 302 to generate a set of intermediate features (not illustrated).” *Id.* at 7:16-

19. The specification provides a “lane-horizon marking head 306” that includes “a CNN configured to process images and output a sequence of points representing the horizon and/or lane lines.” *Id.* at 7:24-27. The specification provides many examples of CNN technologies that may be used with the invention, such as, “a fully convolutional network (FCN) that includes a plurality of convolutional blocks where each block can include one or more convolution, batch normalization, activation and up-sampling layers ... FCN can also be a network comprising, or based on, Unet, DeepLab, PSPNet, SCNN or EnetSAD architectures ... [and] can include a neck module that further processes backbone features to detect markers on fine-grain scale.” *Id.* at 7:27-36. Notably, “the neural network is trained to predict lane and horizon key points 308.” *Id.* at 8:7-8 (emphasis added). The neural network training is based on the processing of images with ground truth image labels to “compute a prediction loss.” *See e.g., id.* at 9:60-10:58 ; see also, FIG. 4A-4B.

22. The specification explains that “the model of FIG. 3A can provide advantages over the offline process of FIG. 2.” *Id.* at 8:37-39. “For example, the use of the neural network can relax constraints that the vehicle needs to be driving on a highway having three or more lane lines. Indeed, just two lane-lines with a known lane-width would be sufficient for camera initialization. That is, since the model is capable of predicting the horizon line directly, previous requirements of

having three lane lines in frame (e.g., to geometrically calculate a horizon line) can be relaxed.” *Id.* at 8:39-46.

23. The inventors described other aspects of the system with reference to “commonly-owned application Ser. No. 17/173,950, filed Feb. 11, 2021” and incorporated this application by “reference in its entirety.” *See e.g., id.* at 5:38-47 (“detail on computing camera properties based on detected lane lines”); 14:34-36 (“detail on computing an IPM”); 15:20-27 (“a geometric algorithm to determine a lane width”) and 49-52)”[d]etails of applying an IPM to a set of image frames to detect lane lines”), and 16:1-11 (“[d]etails of detecting object distances”). Application No. 17/173,950 has since issued as U.S. Patent No. 11,798,187 (“the ’187 patent”). *See* Ex. 2001.

24. The incorporated ’187 patent “claims the benefit of priority of Provisional Pat. Appl. No. 62/975,360 filed on Feb. 12, 2020” and which was itself “incorporated by reference in its entirety” into the ’187 patent disclosure. *See* Ex. 2001 at 1:8-10. Accordingly, the disclosure of Provisional Pat. Appl. No. 62/975,360 is incorporated by reference and therefore part of the disclosure of ’580 Patent. This provisional application includes copies of the Khan (Ex. 1007) and Ali (Ex. 1008) papers relied upon in the Petition.

25. The prosecution history for this patent is quite concise given that in the first office action the Examiner found allowable claims. *See* Ex. 1002 at 0128.

The Examiner noted that claims 4 and 6-7, among others, “would be allowable if rewritten in independent form.” *Id.* at 0136. Of relevance to the Challenged Claims, the Examiner rejected claims 1-3 and 5 under 35 U.S.C. §103 based on KWANT1 et al. (U.S. Pub. No. 2019/0034740 A1) and KWANT2 et al. (U.S. Pub. No. 2019/0102674 A1). *Id.* at 0131; *see also*, Ex. 1057 and Ex. 1058.

26. Because in this Declaration I will focus on the limitation that recites “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line” of claim 1, I will summarize my analysis of the prosecution history with respect to this limitation.

27. For this limitation, in the Office Action the Examiner identified paragraph 46 and the Abstract of KWANT1. Ex. 1002 at 0131. The Examiner quoted from KWANT1’s paragraph 46: “the output of a lane feature detection process include lane models representing the detected lanes or features, and associated prediction confidence values indicating a probability that the processed features of the input image are predictive of the detected lane features.” *Id.* Specifically, as to the lines “including a horizon line,” the Examiner quoted from the Abstract of KWANT1: “determining a horizontal line in the image at which a maximum number of the one or more extended road lanes crosses over a minimum horizontal extent of the horizontal line.” *Id.*

28. The applicants filed a Response that did not amend any claims. *Id.* at 0138-0149. The applicants explained why KWANT1 failed to teach or suggest, among other elements, the “predicting a horizon line” step required by claim 1. *Id.* at 0145. Specifically, the applicants explained:

KWANT1 does not disclose using a predictive model to detect a horizon line. Instead, KWANT1 describes predicting lane lines. KWANT1 then describes extending these lane lines and, using a geometric approach, identifying the horizon line. KWANT1 only describes using predictive modeling to predict lane lines. Thus, KWANT1 does not describe predicting a horizon line as claimed.

Id. at 0145-0146. In response to applicants’ arguments, the Examiner issued a Notice of Allowance allowing all the claims without amendments. *Id.* at 0169.

IV. CHALLENGED CLAIMS

29. The grounds in the Petition seek invalidation of claims 1-7 of the ’580 patent (“Challenged Claims”).

30. Claim 1 is the sole independent claim in this set. Claim 1 is reproduced below:

1. A method comprising:

[a] receiving, over a network from a camera device, a video
comprising a set of image frames;

[b] **identifying one or more lines in the video using a predictive
model, the one or more lines including a horizon line;**

- [c] computing at least one camera parameter based on the one or more lines;
- [d] overlaying the one or more lines on the video to generate an overlaid video;
- [e] transmitting the overlaid video to an annotator device for manual review;
- [f] receiving a confirmation from the annotator device, the confirmation indicating that the one or more lines are accurate; and
- [g] transmitting data representing the at least one camera parameter to the camera device.

31. A key aspect of the Challenged Claims is limitation [b], which during prosecution was highlighted by the inventors as a difference over the prior art identified by the Examiner. As noted above, this Declaration focuses on this limitation and how the Petition has failed to properly show this limitation was known in the prior art.

V. LEGAL STANDARDS

1. CLAIM CONSTRUCTION

32. In rendering the opinions set forth in this declaration, I have read the claims from the perspective of a person of ordinary skill in the art at the time of the invention, in 2021. I understand that in this inter partes review, the claims are

construed in the same manner as in district court, to have their plain and ordinary meaning in the context of the specification and prosecution history.

2. PERSON OF ORDINARY SKILL IN THE ART

33. I understand that a hypothetical person of ordinary skill in the art is considered to have the normal skills and knowledge of a person in a certain technical field, as of the time of the invention at issue. I understand that factors that may be considered in determining the level of ordinary skill in the art include the education level of the inventor; (2) the types of problems encountered in the art; (3) the prior art solutions to those problems; (4) rapidity with which innovations are made; (5) the sophistication of the technology; and (6) the education level of active workers in the field. I also understand that “the person of ordinary skill” is a hypothetical person who is presumed to be aware of the universe of available prior art.

34. The ’580 Patent “relate[s] to automotive systems and to machine learning systems for analyzing sensor data in such systems.” Ex. 1001 at 1:7-10. The type of problem being addressed by the ’580 Patent deals with the initialization of cameras in automotive systems. Moreover, the references relied upon in the Petition are also related to sensor systems in the automotive field. *See e.g.*, Ex. 1005 at [0001] (“disclosure relate generally to operating autonomous vehicles”); Ex. 1007 at 1823 (“[l]ane detection is of critical importance to both

self-driving cars as well as advanced driver assistance systems”); Ex. 1008 (“[d]istance estimation is required for advanced driver assistance systems (ADAS) as well as self-driving cars”); Ex. 1010 at 1:13-14 (“[a] rapidly emerging technology is autonomous vehicles (Avs) that can navigate by themselves ...”). Thus, some experience or training with automotive systems is needed for the appropriate level of ordinary skill in the relevant art.

35. The typical educational level of active workers designing this type of automotive systems in 2021 was at least Bachelor’s Degree in Engineering, typically Electrical, or Computer Science. A typical designer of a system such as those described in the ’580 Patent would have at least two years working experience in the automotive field, perhaps less with a more advanced degree.

36. In my opinion, in October, 2021, a person with ordinary skill in the art with respect to the technology disclosed by the ’580 patent would have at least a Bachelor’s Degree in Engineering or Computer Science with one or two years training and/or experience in design of software for camera/sensor systems in the automotive industry.

37. Petitioner’s proposed POSITA lacks any training or experience in “automotive systems” which is the main technology area of the ’580 patent. In the automotive industry, the type of problems a POSITA would encounter are not the same type of problems that someone doing research or working in more generic

computer vision technology. In the automotive industry, the problems that arise typically involve the application of other technologies to solve problems specific to vehicles, driving, etc. While a person with the skills and background the Petitioner proposes may be more closely working in research or development of underlying computer vision algorithms or technology. Thus, in my opinion, Petitioner's proposed POSITA is not correct with respect to the '580 patent.

38. In sum, based on my experience, training, and education, I consider myself (both now and as of October 2021) to be a person of at least ordinary, or expert, skill in the art with respect to the field of technology implicated by the '580 patent.

3. OBVIOUSNESS

39. I understand that a patent claim is unpatentable and invalid if the subject matter of the claim as a whole would have been obvious to a person of ordinary skill in the art of the claimed subject matter as of the time of the invention at issue. I understand that the following factors must be evaluated to determine whether the claimed subject matter is obvious: (1) the scope and content of the prior art; (2) the difference or differences, if any, between each claim of the patent and the prior art; and (3) the level of ordinary skill in the art at the time the patent was filed. Unlike anticipation, which allows consideration of only one item of prior art, I understand that obviousness may be shown by considering more than one

item of prior art. Moreover, I have been informed and I understand that so-called objective indicia of non-obviousness, also known as “secondary considerations,” like the following are also to be considered when assessing obviousness: (1) commercial success; (2) long-felt but unresolved needs; (3) copying of the invention by others in the field; (4) initial expressions of disbelief by experts in the field; (5) failure of others to solve the problem that the inventor solved; and (6) unexpected results. I also understand that evidence of objective indicia of non-obviousness must be commensurate in scope with the claimed subject matter.

VI. ANALYSIS OF GROUNDS IN THE PETITION

1. Ground 1

Choe (Ex. 1005)	Claims 1-7	Section 103
-----------------	------------	-------------

40. Petitioner asserts that Choe (Ex. 1005) renders obvious the Challenged Claims. It does not. At the very least, Choe fails to teach or suggest limitation [b] of claim 1, “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line.”

a) Horizon Line Determination in Choe

41. Choe discloses a manual camera calibration system that uses camera hardware to determine the horizon line. As the specification describes, “a

calibration system with a visualizer is utilized to calibrate sensors (e.g., cameras, LIDAR, RADAR) of an autonomous driving vehicle (ADV).” Ex. 1005 at [0014]. “The visualizer includes a user interface to visually display a two-dimensional (2D) image captured by a camera and a three-dimensional (3D) image to allow a user visually calibrate the camera, a LIDAR device, and/or a RADAR device.” *Id.* “A user can calibrate any of the sensors such as a pitch angle, a yaw angle, and/or a roll angle using an input device such as a keyboard or joystick to provide incremental calibration feedback, while displaying the calibration result on a display device concurrently.” *Id.*

42. Choe describes that “in response to a first image captured by a camera of an ADV, a horizon line is determined based on the camera's hardware settings.” *Id.* at [0015]. “The horizon line represents a vanishing point from a view point of the ADV based on an initial or default pitch angle of the camera.” *Id.* “One or more lane lines are determined based on the first image via a perception process performed on the first image.” *Id.* “In response to a first input signal (e.g., an up or down arrow key) received from an input device (e.g., keyboard, joystick, voice interactive command), a position of the horizon line is updated based on the first input signal and a position of at least one of the lane lines is updated based on the updated horizon line.” *Id.*

43. Choe consistently describes throughout the specification that the horizon line is determined based on the hardware parameters of the camera. *See e.g.*, Ex. 1005 at [0046] (“2D rendering module 401 is configured to render a 2D view including superimposing a horizon line onto the image according to the camera setting (e.g., hardware preset pitch angle).”); at [0048] (“perception module 302 determines a horizon line based on the camera’s hardware settings.”); at [0055] (“horizon line 610 is determined and placed on a location within image 601 based on the hardware setting of the targeted camera.”); at [0059] (“processing logic determines a horizon line ... based on a pitch angle of a camera that captured a first image”).

44. Choe describes a “machine learning engine 122,” that “[b]ased on driving statistics 123, ... generates or trains a set of rules, algorithms, and/or predictive models 124 for a variety of purposes.” *Id.* at [0030]. Choe states that “[a]lgorithms 124 may further include sensor calibration algorithms, which will be described further below.” *Id.* Choe further describes the use of “sensor calibration algorithms” as follows:

[S]ensor calibration module or system 308 is configured to calibrate, using sensor calibration algorithms 313, a sensor based on an image or point cloud captured by the sensor to determine a set of calibration parameters. The calibration parameters can be utilized by the software to compensate the hardware settings of the sensors during a perception process. Sensor calibration module 308 may be implemented as a part of perception module 302 or communicatively coupled to perception module 302. The functionalities of sensor calibration module 308 may

be maintained in a data analytics system such as server 103 to perform a sensor calibration of a sensor offline based on the images captured by the corresponding sensor. The calibration result can then be uploaded onto the vehicle to be utilized online during the image processing as a part of the perception process.

Id. at [0044]. But as I explain below, these sensor calibration algorithms are not predictive models at all.

45. Choe describes that “a data analytics system to perform data analytics services for a variety of clients ... and includes data collector 121 and machine learning engine 122.” *Id.* at [0029]. “Data collector 121 collects driving statistics 123 from a variety of vehicles, either autonomous vehicles or regular vehicles driven by human drivers.” *Id.* “Driving statistics 123 include information indicating the driving commands (e.g., throttle, brake, steering commands) issued and responses of the vehicles (e.g., speeds, accelerations, decelerations, directions) captured by sensors of the vehicles at different points in time.” *Id.* “Driving statistics 123 may further include information describing the driving environments at different points in time, such as, for example, routes (including starting and destination locations), MPOis, road conditions, weather conditions, etc.” *Id.*

46. The “set of rules, algorithms, and/or predictive models 124” of Choe are generated by a machine learning engine 122 “based on driving statistics 123.” *Id.* at [0044]. Thus, these “set of rules, algorithms, and/or predictive models 124” are based on “driving statistics 123.” Algorithms 124 are used for sensor

calibration “to determine a set of calibration parameters.” Nothing in Choe teaches or suggests “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line.”

b) Choe Does Not Render Claims 1-7 Obvious

47. The Petition states that in Choe, “[t]he initial horizon and lane lines are identified ‘using a predictive model.’” Pet. at 27. This is incorrect. Choe consistently describes throughout the specification that the horizon line is determined based on the hardware parameters of the camera, not using a predictive model. *See e.g.*, Ex. 1005 at [0046], [0048], [0055], [0059].

48. The Petition relies on Choe’s paragraph [0030] referring to “predictive models 124” possibly used for sensor calibration. Pet. at 27 (citing Ex. 1005 at [0030]). From this disclosure, Petitioner incorrectly concludes that “the initial horizon line is predicted with a predictive model 124.” *Id.* at 27-28. But Paragraph [0030] in Choe states that “[b]ased on driving statistics 123, machine learning engine 122 generates or trains a set of rules, algorithms, and/or predictive models 124 for a variety of purposes.” Ex. 1005 at [0030]. That is, any “predictive models” in Choe are based on the use of machine learning module to process driving statistics, such as vehicle speeds, accelerations, etc., not based training a predictive model with images to identify lines in the images.

49. Additionally, Paragraph [0030] in Choe clearly refers to its additional description of “sensor calibration.” Ex. 1005 at [0030] (noting use of algorithms for sensor calibration “will be described further below”). But the Petition and the corresponding expert declaration do not analyze this additional description. What Choe describes is its “sensor” calibration approach. Sensor calibration in Choe includes not only camera but also LIDAR devices, RADAR devices; Choe is not specific to “camera” calibration. Further, calibration itself is a broad term and does not necessarily include identifying a horizon line.

50. With respect to sensor calibration, Choe describes that a “sensor calibration module or system 308 is configured to calibrate, using sensor calibration algorithms 313, a sensor based on an image or point cloud captured by the sensor to determine a set of calibration parameters.” Ex. 1005 at [0044]. Thus, the “algorithms 124” referenced in paragraph [0030], which may be based on machine learning processing of driving statistics (e.g., driving commands, speeds, accelerations, decelerations, etc.), are used for sensor calibration “to determine a set of calibration parameters,” not a horizon line. *See id.* at [0030] and [0044].

51. Th predictive models in Choe are also fundamentally different from the predictive models of the '580 patent because they are not trained on “ground truth image labels corresponding to the processed image along with all the classification prediction head's output.” *See e.g., id.* at 9:60-10:58 ; see also, FIG.

4A-4B. Instead, the Choe models, at best, may be trained on vehicle motion driving statistics. Ex. 1005 at [0030]. This reliance on driving-based statistics creates at least two core limitations when it comes to identifying horizon lines in general video data. First, these models implicitly assume that the camera is rigidly mounted on a vehicle following known dynamics (e.g., forward motion along drivable surfaces), and their training data reflects only those conditions. As such, the models cannot generalize well to video captured from different perspectives, camera placements, or vehicle behaviors—especially in cases of dynamic pitch, unmodeled road curvature, or off-road environments.

52. Second, because any training process in Choe would have to use motion-based statistics rather than labeled images, such models could never explicitly learn to interpret complex visual indicators of the horizon—such as cloud lines, skyline silhouettes, or vanishing points in unstructured scenes. They can only be based on calibration parameters that correlate with the specific geometry of the vehicle’s past motion statistics.

53. Thus, Choe does not teach or suggest the use of a predictive model for any purpose, much less for determining a horizon line in video images. The only teaching regarding predictive models in Choe is that they are “based on driving statistics” and “generated or trained” by “machine learning engine 122.” Nothing in Choe teaches or suggests any predictive model trained to identify lines in video.

Therefore, the driving-statistics based models of Choe are fundamentally unsuitable for generalized horizon detection tasks.

54. But more importantly, to determine the horizon line, Choe teaches the use of hardware parameters of the camera, not the use of predictive models. Thus, Choe does not teach or suggest using predictive models for “identifying one or more lines in the video ... including a horizon line” as required by claim 1.

2. Ground 2

Choe (Ex. 1005) Khan (Ex. 1007) Ali (Ex. 1008)	Claims 1-5	Section 103
--	------------	-------------

55. Ground 2 combines Khan and Ali to Choe, in order to further show the “predictive model” limitations of the claims but also fails to show limitation [b].

a) Horizon Line Determination in Khan and Ali

56. Kahn describes lane detection techniques for self-driving and advanced driver assistance systems. Ex. 1007 at Abstract. The authors describe their “use [of] a lane boundary marker network to detect keypoints along the lane boundaries.” Ex. 1007 at Abstract. The authors note that instead of using “deep learning methods for lane detection that output lane boundary maps which are sampled to produce boundary markers,” they take “motivation from keypoint

detection in Mask-RCNN [12] and predict these lane markers directly instead of sampling them from segmentation maps.” *Id.* at 1825.

57. Section “3.2.1 Horizon Estimation” describes that “[t]he estimation of IPM assumes a planar road scene having no elevation and more than one lanes (3 or more lane boundaries), such as a typical highway scenario.” *Id.* at 1826. Kahn describes that a “[c]ritical part of initialization is the estimation of a vanishing line (horizon) from these lines.” *Id.* at 1826. Kahn’s methodology for estimating the horizon “requires two sets of vanishing points (VPs), forward and lateral VPs.” *Id.* “Forward VPs can be found from the cross product of the world parallel lines whereas the process of finding lateral vanishing point is a bit involved.” *Id.* at 1827. “For that we use cross-ratios of three real-world parallel lines which are the lane lines as used by Ali et al. [2].” *Id.* (citing the Ali paper of Ex. 1008). “Once a set of forward and lateral VPs is found, we use their horizon estimation algorithm to find the best fitting horizon l_h .” *Id.*

58. Thus, Khan does not disclose a specific approach to determine the horizon line but rather refers to the approach disclosed in Ali (Ex. 1008) to estimate the horizon.

59. Ali describes distance estimation techniques for advanced driver assistance systems (ADAS) and for self-driving cars. Ex. 1008 at Abstract. Ali states that it “introduces novelty on two fronts: (1) it uses cross-ratios of lane

boundaries to estimate horizon [and] (2) it determines an Inverse Perspective Mapping (IPM) and camera height from a known lane width and the detected horizon.” Ex. 1008 at Abstract.

60. Section 3.2.2 (“Horizon estimation”) describes that “[t]he horizon must be known to estimate the required camera rotation for an IPM.” *Id.* at 1103. For the Ali system, “[a] horizon is a vanishing line of a plane which, in our case, is the road.” *Id.* Specifically, Ali explains:

In order to find vanishing line, we need two vanishing points. The first one can be simply calculated through the cross product of the vectors of any two lane lines expressed in homogeneous coordinates (see Figure 5 (a)). We call it the forward vanishing point (V). Finding a second vanishing point, however, is a bit involved. ... For this, we will need three co-linear points existing on the road plane having known real world length ratios. Since we already have detected lane lines in the image which lie on the road plane, we can find three co-linear points by drawing any line l' which intersects all three lane lines at points a' , b' and c' respectively.

Id. at 1103-1104. The process of getting the vanishing points is illustrated in Figure 5, which is reproduced below:




Figure 5. (a) Perspective view of road plane. Lane lines (l'_1 , l'_2 , and l'_3) are shown in red and an arbitrary line l' (blue) cuts all lane lines at points a' , b' , and c' . Forward vanishing point (V) is obtained by the cross product of any two lane lines (b) Real world top-down view of the road plane. d_1 and d_2 are the distances between lane lines l_1 , l_2 and l_2 , l_3 respectively. Since all lane lines are parallel Δ_1 and Δ_2 are similar triangles which implies $\frac{ab}{bc} = \frac{d_1}{d_2}$

Id. at 1104. The fact that the three VPs “lie on three (parallel) lane lines (l_1 , l_2 and l_3), it can be established through trigonometry that the ratio of the line segments ab/bc is equal to d_1/d_2 where d_1 and d_2 are the distances between lines l_1 , l_2 and l_3 respectively.” *Id.*

61. Ali describes that it is helpful to know that initialization may be done on a highway or section of a highway where lanes are consistent/equidistant (using

gps)” because this allows to assume that the ratio of the distances between the lines to be 1. *Id.* at 1104. “After getting two vanishing points (VPs) per frame,” Ali describes that the “horizon” can be found using “all the sets of the vanishing points across all the frames of a video clip” by finding “the best fitted line to the VPs as described in Algorithm 1.” Ex. 1008 at 1104. This is the “[b]est fitting Horizon (l_h) estimation” algorithm referred to in the Kahn paper. *See* Ex. 1007 at 1827

62. Thus, Ali describes a purely geometrical approach to determining the horizon.

b) b. The Choe-Khan-Ali Combination Does not Render Claims 1-7 Obvious

63. The disclosure of the ’580 patent includes the same disclosure provided by Khan and Ali. Khan and Ali are part of the disclosure of the ’580 Patent through incorporation by reference. The specification refers to the disclosure of Khan and Ali as providing “additional detail” regarding some aspects of the invention. However, Khan and Ali do not disclose limitation [b] of claim 1, “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line.” Thus, Khan and Ali do not make up the missing disclosure of Choe discussed above with reference to Ground 1.

64. While Petitioners allege that “[i]n Khan, the predictive model used to identify lane and horizon lines is called a ‘Mask-RCNN’,” there is no such disclosure in the paper. *See* Pet. at 41 (citing Ex. 1007, 1825). Instead, the paper

describes that “[a]s generally is the case with deep learning methods for lane detection that output lane boundary maps which are sampled to produce boundary markers, we instead take motivation from keypoint detection in Mask-RCNN [12] and predict these lane markers directly instead of sampling them from segmentation maps.” Ex. 1007 at 1825. Petitioner is correct that the inventors did not invent “Mask-RCNN” (Pet. at 15), since the authors clearly referred to item no. 12 in their bibliography as the source of their “motivation.” *Id.* The cited paper, “K. He, G. Gkioxari, P. Dollár, and R. B. Girshick. Mask RCNN. International Conference on Computer Vision (ICCV), pages 2980–2988, 2017” appears to be Ex. 1009 to the Petition. In this paper, the authors “present a conceptually simple, flexible, and general framework for object instance segmentation.” Ex. 1009 at 1. This shows that the Mask-RCNN model is a generic framework that can be used for different computer vision applications. The fact that the disclosure of the ’580 patent, including the Khan and Ali papers, all refer to the use of this framework in the context of automotive sensor systems it not surprising since it is “simple” and “flexible.” However, contrary to what the Petition states, what Khan does not describe anywhere in the paper is using the flexible Mask-RCNN framework to identify a horizon line in video images.

65. Petitioner also incorrectly states that “Khan uses RANSAC (a predictive model) to predict the horizon line.” *See* Pet. at 41 (citing Ex. 1007,

1816-27). First, RANSAC is not a predictive model, or any sort of model at all. RANSAC, which stands for “Random Sample Consensus,” provides a process for fitting a model to experimental data finding a match that has the least number of outliers. What Khan actually describes in this context is the following:

3.2.1 Horizon Estimation

The estimation of IPM assumes a planar road scene having no elevation and more than one lanes (3 or more lane boundaries), such as a typical highway scenario. We start by fitting lines to the detected markers. A best fit line must have a minimum number of inlier markers (i.e., 25% of total number of markers for each lane boundary). Inlier markers are those having horizontal distance of under 8 pixels to the best 1st degree line fit through RANSAC [8].

Ex. 1007 at 1826. I noted that for the computation of IPM, the ’580 patent explicitly refers to and incorporates this disclosure as providing further details. *See* Ex. 1001 at 14:34-36. Thus, in Khan, RANSAC, bibliographic reference no. 8 to the Khan paper, is referenced in the context of defining an “inlier marker.” But contrary to the Petitioner’s statement, Khan does not describe using RANSAC to predict the horizon line.

66. Khan references Ali (Ex. 1008), which is bibliographic reference no. 2 to the Khan paper, for the use of “cross-ratios of three real-world parallel lines which are lane lines” and Khan’s use of Ali’s “horizon estimation algorithm to find the best fitting horizon l_h .” Ex. 1007 at 1827.

67. Ali describes that it needs “three co-linear points existing on the road plane having known real world length ratios,” which they find by “drawing any line l' which intersects all three lane lines,” to determine a horizon. Ex. 1009 at 1104-1104. As I explained above with reference to Ali’s Fig. 5, Ali describes that since three VPs “lie on three (parallel) lane lines (l_1 , l_2 and l_3), it can be established through trigonometry that the ratio of the line segments ab/bc is equal to d_1/d_2 where d_1 and d_2 are the distances between lines l_1 , l_2 and l_3 respectively.” Ex. 1008 at 1104 and Fig. 5.

68. Since Khan uses the Ali approach, they also need to “assume[] a planar road scene having no elevation and more than one lanes (3 or more lane boundaries), such as a typical highway scenario.” Ex. 1007 at 1826.

69. The ’580 patent specification expressly noted that the use of predictive models for identifying the horizon line improved over these types of geometric approaches. As the ’580 Patent explains:

the use of the neural network can relax constraints that the vehicle needs to be driving on a highway having three or more lane lines. Indeed, just two lane-lines with a known lane-width would be sufficient for camera initialization. That is, since the model is capable of predicting the horizon line directly, previous requirements of having three lane lines in frame (e.g., to geometrically calculate a horizon line) can be relaxed.

Ex. 1001 at 8:39-46. This geometric calculation of the horizon line is the type of geometric approach for estimating the horizon described by Ali and used by Khan,

which imposes constraints that “can be relaxed” with the use of the ’580 patent “predictive model” approach.

70. Thus, the Khan/Ali geometric horizon estimation disclosure does not teach or suggest limitation [b] of claim 1, “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line.” Moreover, Choe repeatedly describes that “a horizon line is determined based on the camera’s hardware settings.” Ex. 1005 at [0015]. Absent the teachings of the ’580 patent, a POSITA would not be motivated to replace Choe’s hardware-based approach to determine horizon lines with Khan/Ali geometric approach to identify a horizon line. But the combination of Khan/Ali’s geometric approach to determining a horizon line with the Choe system, even if it were appropriate and not based on the ’580 patent’s disclosure, still does not show the Challenged Claims to be obvious because, it does not disclose the claimed “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line” since, as I explained above with reference to Ground 1, Choe also fails to teach or suggest this same limitation.

3. Ground 3

Choe (Ex. 1005) Workman (Ex. 1011)	Claims 1-4	Section 103
---------------------------------------	------------	-------------

71. Ground 3 of the Petition is based on Ground 1 as to its analysis of Choe and combines with it Workman (Ex. 1011) for its generic computer vision disclosures. Workman describes research in computer vision “investigat[ing] methods for directly estimating the horizon line using convolutional neural networks.” Ex. 1011 at 10. The paper notes that “[t]he horizon line is an important contextual attribute for a wide variety of image understanding tasks.” *Id.* at Abstract. However, Workman does not mention any automotive systems or initialization of dash-cam or automotive sensor systems as a relevant field or application of its experimental approach. The paper “introduce[s] a new benchmark dataset” called “*Horizon Lines in the Wild* (HLW) containing natural images with labeled horizon lines.” *Id.* The authors found this new dataset needed to improve over pre-existing datasets for training neural networks to identify horizon lines, because existing datasets were deemed insufficient. *Id.* at 1. It also describes that “[d]espite [the] demonstrated importance [of horizon line estimation in computer vision], progress on this task has stagnated and nearly all recent methods that focus on this problem make assumptions about the presence of particular geometric objects in the scene, such as vanishing points, repeated textures, and coplanar circles.” *Id.* (citations omitted). In fact, Workman notes that “[r]ecent work using learning-based methods for horizon line estimation has been limited.” Ex. 1011 at 2. In this context, the authors “investigate the

application of convolutional neural networks for directly estimating the horizon line, without requiring any explicit geometric constraints or other special cues.” *Id.* at Abstract. They used a CNN trained with their newly introduced HLW data set combined with street-side imagery. *Id.* at 8.

72. Thus, in my opinion, Workman is not applicable prior art that a POSITA would have looked to in addressing automotive system problems in the prior art, whether under Petitioner’s definition of a POSA or my own definition. Particularly, under my definition of a POSITA, which I believe is the correct one, someone working in design of automotive sensor systems would not be motivated to combine an experimental computer vision approach for horizon line determination to the Choe system. Workman is a generic computer vision reference that has no relation to the type of automotive system of Choe.

73. Further, as explained above, Choe uses a camera-hardware-based approach to determine initial horizon lines that then are manually adjusted. Choe repeatedly describes that “a horizon line is determined based on the camera’s hardware settings.” *Id.* at [0015]. Absent the teachings of the ’580 patent, a POSITA would not be motivated to replace Choe’s hardware-based approach to determine horizon lines with Workman’s image-trained predictive models to identify a horizon line. This is particularly the case given that Choe discloses using predictive models, not trained with images but rather based on driving

statistics, like speeds, accelerations, etc., and teaches the use of algorithms 124 not to determine the horizon line but for sensor calibration using “a set of calibration parameters.” *See* Ex. 1005 at [0029], [0030], and [0044].

74. Thus, in my opinion, Choe’s teaching of the use of driving-statistics-based predictive models, even if they were used to calibrate sensors with calibration parameters, as Petitioner alleges, would not motivate a POSITA to look for other references, much less references disclosing the use of image-trained predictive models to identify horizon lines. There is simply nothing in Choe to motivate a POSITA to look for other references discussing the use of image-trained predictive models, much less for the different purpose of identifying a horizon line. Choe does not identify any shortcomings with its camera-hardware based approach to determining the horizon line or the need for using predictive models for this purpose. Nor does Choe identify or suggest any problem or other need to improve its sensor calibration approach in any way.

75. In addition, the statements in Workman that Petitioner points to as supporting this alleged motivation to combine do not actually support the proposed combination with Choe.

76. With respect to the “state-of-the-art” performance, the Workman paper results do not relate to any comparison with initialization of cameras in vehicular systems or in any sort of automotive system. The “state-of-the-art”

performance referenced in the Petition is that of the described experimental system “on the challenging HLW dataset and two existing benchmark datasets.” Ex. 1011 at 1. The fact that the experimental system described in Workman matches existing “state-of-the art” approaches on benchmark datasets provides no indication of whether a POSITA would use Workman’s predictive models to identify the horizon lines in Choe, or then subsequently use that horizon line as the basis for calibration of the camera. As noted above, Choe does not use any sort of predictive model to identify horizon lines; it uses the camera’s hardware settings. The alleged “algorithms/models 124” of Choe, which as Petitioner interprets Choe, are not used for identifying lines but to generate “calibration parameters,” are based on the machine learning processing of “driving statistics” and have no relation or ability to work with the HLW dataset of “real-world images with labeled horizon lines” or any other dataset of images.

77. In short, Choe uses camera hardware settings to determine the horizon line and there is no indication that its results are not already “state-of-the-art” without the need for the combination with Workman. A POSITA would not know how or why the Workman disclosure is relevant to the completely different Choe system.

78. With respect to the “faster” performance, the “orders of magnitude faster” reported in Workman is not relevant to Choe. Instead, the authors

compared their proposed experimental approach to “current state of the art methods” but do not mention any methods used in automotive systems, much less calibration of automotive sensors. Ex. 1011 at 1. Workman describes that “[r]ecent state-of-the-art methods are quite slow, which is reasonable when you have a small testing dataset.” *Id.* at 3. For example, they “find that the approach of Lezama et al. [19] requires approximately 30 seconds per image on YUD and 1 minute per image on ECD (results obtained using code made available by the authors).” *Id.* The fact that the experiments in Workman required “only milliseconds per image” and “outperform[ed] all existing methods” does not provide any indication of how it performs compared to the camera-hardware-based approach of Choe. *See* Pet. at 48.

79. Choe does not report any speed performance issues with this approach. In fact, in the Choe calibration process, the horizon line is initially determined by on the camera hardware settings, which is likely a very fast approach, and then it allows for a manual adjustment from the user: “a position of the horizon line is updated based on an input signal” from the user, “received from an input device (e.g., keyboard, joystick, voice interactive command.)” *See* Ex. 1005 at [0015]. Given this user-based refinement of the horizon line position, the difference between 30 seconds or milliseconds per image may not be that relevant to the Choe system. Thus, this reported faster performance in Workman does not

show that a POSITA would use Workman’s image-trained predictive models to replace the hardware-based identification of the horizon lines in Choe, a completely different approach.

80. In sum, in my opinion, the Petition does not identify any motivation in Choe for a POSITA to go look for other publications disclosing the use of image-trained predictive models to identify horizon lines in video images. Choe discloses the use of predictive models based on driving statistics for sensor calibration and self-driving assistance but does not suggest their use for horizon line identification. Workman is in an unrelated field as it does not pertain to automotive systems and provides no motivation to combine its image-trained models with a sensor calibration system like Choe’s. Neither the Petitioner nor its expert provide any rationale for why a POSITA would have been motivated to change the “camera-hardware-based” approach for identifying the horizon line in Choe with the type of image-trained algorithms disclosed in Workman.

4. Ground 4

Choe (Ex. 1005) Westmacot (Ex. 1010).	Claims 5-7	Section 103
--	------------	-------------

81. Westmacot “relates to image annotation, such as the creation of annotated street road images which can be used (among other things) for training a machine learning road detection component.” Ex. 1010 at 1:5-7.

82. Because the Petition does not rely on Westmacot with respect to the “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line” limitation, I do not summarize my review of Westmacot at this time.

83. Because Ground 4 relies on the analysis of Choe in Ground 1 with respect to claim 1, as already explained above, Ground 4 also fails as to dependent claims 5-7. Specifically, Petitioners do not identify any additional disclosure in Westmacot that would teach or suggest the “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line” element required by claim 1 and missing in Choe.

5. Ground 5

Choe (Ex. 1005) Khan (Ex. 1007) Ali (Ex. 1008) Westmacot (Ex. 1010).	Claims 1-7	Section 103
---	------------	-------------

84. Ground 5 relies on Petitioner’s arguments regarding Grounds 1, 2, and 4. Because, as I explained above with respect to each of those grounds, neither

Choe, Westmacot, Khan, or Ali teach or suggest the “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line” element required by claim 1, Petitioner’s analysis of Ground 5 suffers of the same deficiencies.

6. Ground 6


Choe (Ex. 1005) Westmacot (Ex. 1010) Workman (Ex. 1011)	Claims 1-7	Section 103
---	------------	-------------

85. Ground 6 relies on Petitioner’s arguments regarding Grounds 1, 2, and 3. Accordingly, my discussions of those grounds apply to Ground 6. As explained above with respect to grounds 1 and 2, neither Choe, nor Westmacot teach or suggest the “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line” element required by claim 1. Further, as explained with respect to Ground 3, the combination with Workman is improper, lacking motivation, and based on the ’580 patent disclosure.

VII. ADDITIONAL ANALYSIS OF GROUNDS 1, 3, AND 4 IN THE PETITION

A. Ground 1

86. Choe does not disclose “receiving, over a network from a camera device, a video comprising a set of image frames.” The Petition points to Figure 6A and 6B from Choe to show “video frames” in Choe’s manual calibration process. Pet. at 25. The Petition then asserts that “[t]he **video** captured by the video camera—a sensor of sensor system 115—is received, over the network, by the sensor calibration system 125 on server(s) 103. EX1005, ¶[0031], FIG. 1.” *Id.* (emphasis added). However, neither paragraph 31 nor FIG. 1 of Choe refer to any video, much less receiving video over a network from a camera device. Instead, Choe shows its operation is based on single frames. For example, FIG. 7 shows this operation:

**FIG. 7**

Ex. 1005, Fig. 7. The initial steps of the process, 701 and 702, refer to a/the “first image.” This single-image operation is further reflected by the fact that Choe uses multiple sensors , including “one or more cameras 211, global positioning system (GPS) unit 212, inertial measurement unit (IMU) 213, radar unit 214, and a light detection and range (LIDAR) unit 215.” *Id.* at [0023]. In fact, the manual calibration process that the Petition relies upon, illustrated in FIG. 6A and 6B, uses “other 3D information, for example, provided by a LIDAR device and/or RADAR device” in addition to the camera image. *Id.* at [0056]. This requires a tight correlation between each image and the data provided by other sensors (e.g., LIDAR/RADAR), which is simplified when operating on a single image at a time,

instead of on a running video. Thus, a POSITA would not understand Choe to disclose any operation on a video, but rather on single images. Thus, Choe does not disclose “receiving, over a network from a camera device, a video comprising a set of image frames.”

87. In addition to the reasons set forth above, Choe does not disclose the predictive models of element [1B] for additional reasons. As I noted in paragraph 48, the Petition relies on Choe’s paragraph [0030] referring to “predictive models 124” possibly used for sensor calibration. Pet. at 27 (citing Ex. 1005 at [0030]). As I discussed above, at most, Choe describes that the “predictive models 124” referenced in paragraph [0030], which are based on machine learning processing of driving statistics (e.g., driving commands, speeds, accelerations, decelerations, etc.), are used for sensor calibration “to determine a set of calibration parameters,” not a horizon line. *See id.* at [0030] and [0044]. But in fact, Choe does not actually describe using “predictive models” even for camera calibration.

88. Choe refers generically to “a set of rules, algorithms, and/or predictive models 124” all together as being generated or trained by machine learning engine 122 “for a variety of purposes.” *Id.* at [0030]. Petitioner and Dr. Darrell conclude from this statement that “predictive models” are used for sensor calibration. Pet. at 27, EX 1003 at ¶¶93-94.

89. But Choe does not refer to “sensor calibration algorithms/models” like Petitioners do. Instead, in the generalized statement of paragraph [0030], Choe refers to “rules,” “algorithms,” and/or “predictive models” all of which are referred to as element 124 being “generated” or “trained” by “machine learning engine 122.” This makes a clear distinction between “algorithms” and “predictive models” as being different in Choe. In fact, a POSITA would not understand what it means to “train” an algorithm. The next reference to element 124 in Choe is that “Algorithms 124 may further include sensor calibration algorithms, which will be described further below.” Thus, it is algorithms 124, not predictive models, that are used for sensor calibration in Choe.

90. Choe then explains how the “set of rules” are used. For example, “Decision module 304 may make such decisions according to a set of rules such as traffic rules or driving rules 312, which may be stored in persistent storage device 352.” EX 1004 at [0038]. It also explains how “algorithms” are used for “autonomous driving a vehicle,” for “sensor calibration,” and how they can be uploaded and “utilized during autonomous driving in real-time.” *Id.* at [0030]. Choe does not actually describe how “predictive models” are used for any purpose, how they may be trained, or provide any specifics about them. More importantly, nothing in Choe teaches or suggests the use of a predictive model to identify a horizon line.

91. As I already noted above, specifically as to calibration, Choe describes that “sensor calibration module or system 308 is configured to calibrate, using sensor calibration **algorithms** 313, a sensor based on an image or point cloud captured by the sensor to determine a set of calibration parameters.” *Id.* at [0044]. Thus, it is the “algorithms” generated by “machine learning engine 122” that are used for sensor calibration in Choe, not the “predictive models” or the “set of rules.” Thus, in Choe

92. Every paragraph cited in the Petition alleging the disclosure of a predictive model used to predict camera parameters with respect to claim 3 describes instead a manual adjustment process based on user input. Pet. at 35. The Petition states: “These sensor calibration algorithms/models predict initial camera parameters and updated camera parameters. EX1005, ¶¶[0015], [0017], [0047], [0049], [0052]-[0054], [0059].” However, none of these paragraphs show predictive models used to identify initial or updated camera parameters:

- [0015]: **In response to a first input signal (e.g., an up or down arrow key) received from an input device (e.g., keyboard, joystick, voice interactive command), a position of the horizon line is updated based on the first input signal and a position of at least one of the lane lines is updated based on the updated horizon line. The input signal may represent an incremental adjustment for adjusting the**

position of the horizon line. A first calibration factor or first correction value is determined for calibrating a pitch angle of the camera **based on a difference between the initial horizon line and the updated horizon line.**

- [0017]: According to one embodiment, **in response to a second input signal received from the input device (e.g., left or right arrow key),** the position of at least one of the lane lines is modified based on the second input signal. The positions of the lane lines are updated on the second image displayed on the second display area. A second calibration factor for calibrating a yaw angle of the camera **is determined based on the modification of the positions of the lane lines.** Similarly, **in response to a third signal received from the input device (e.g., shift-left or shift-right arrow key),** the lane lines are modified and a third calibration factor for calibrating a roll angle of the camera is determined. The first, second, and third images are concurrently displayed within the first, second, and third display areas respectively. The processes similar to the above process can be implemented to determine calibration factors for LIDAR and/or RADAR devices. ...

- [0047]: A user can then enter feedback or adjustment **using input device 404 such as a keyboard or joystick. In response** to the input provided by the user, modules 411-413 **adjust certain parameters**, render the images again, and display the updated images on display device 405. As a result, the user can calibrate the sensors until the calibration results as shown on display device 405 are satisfied.
- [0049]: **In response to a first input signal (e.g., an up or down arrow keypress) received from an input device (e.g., a keyboard, a joystick, a microphone to receive voice interactive commands),** a position of the horizon line is updated based on the first input signal and a position of at least one of the lane lines is updated based on the updated horizon line. The input signal may represent an incremental value for adjusting the position of the horizon line. A first calibration factor or first correction value **is determined** for calibrating a pitch angle of the camera **based on a difference between the initial horizon line and the updated horizon line.**
- [00552-54]: According to one embodiment, **in response to a second input signal received from the input device 404 (e.g., left or right arrow key),** the position of at least one of the lane lines is modified based on the second input signal. The positions of the lane lines are

updated on the second image 412 displayed on the second display area. A second calibration factor for calibrating a yaw angle of the camera **is determined based on the modification of the positions** of the lane lines. ... Similarly, **in response to a third signal received from the input device (e.g., shift-left or shift-right arrow key)**, the lane lines are modified and a third calibration factor for calibrating a roll angle of the camera **is determined**. ...

- [0059]: **In response to a first input signal received from an input device**, at block 703, the position of the horizon line is updated based on the first input signal. At block 704, a position of at least one of the lane lines is updated based on the updated horizon line. At block 705, processing logic **determines a calibration factor or parameter for calibrating a pitch angle of the camera based on the difference between the initial horizon line and the updated horizon line**.

93. In short, none of the paragraphs identified in the Petition with respect to Claim 3 shows any predictive model being used to predict any camera parameters; the references to the specification all show the use of manual user input to adjust the lane lines and using the manually entered difference to determine calibration factors.

B. Ground 2

94. Claim 2 requires that “the one or more lines includes at least one lane line of a roadway.” Ex. 1001, claim 2. This reference to “the one or more lines” makes clear that these are the lines identified using a predictive model:

“identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line.” Ex. 1001, claim 1. Further, given that limitation [1B] requires the using the predictive model to identify the one or more lines, which must include a horizon line, the same predictive model must be shown to being used to identify the horizon line of limitation [1B] and the “lane line of a roadway” of claim 2.

95. I have reviewed the allegations in the Petition regarding claim 2 and in my opinion they fail to show the same predictive model being used to identify a horizon line and a lane line of a roadway. *See* Pet. at 41-43. A POSITA would not consider the “end2end detection algorithm” of Figure 3 in Khan (Ex. 1007), to be a predictive model. Khan refers to this algorithm as including a “post-processing pipeline” that follows “lane boundary markers detection.” Ex. 1007, p.1825. The Khan approach includes a “lane boundary keypoint detection” process that is different than the “camera initialization” process, part of which includes a “horizon estimation.” *Id.* at 1826. In the horizon estimation process, Khan refers to Ali’s “horizon estimation algorithm.” *Id.* at 1827. However, Khan does not rely on

Ali's horizon line estimation approach for identifying lane lines of a roadway in the video.

C. Ground 3

96. Workman is about using a CNN for single-image horizon estimation in natural photographs; it says nothing about video pipelines, networking, overlay generation, annotator confirmation, or sending parameters back to the camera. *See* Ex. 1011.

97. Choe describes that “in response to a first image captured by a camera of an ADV, a horizon line is determined based on the camera's hardware settings.” Ex. 1005, ¶[0015]. “The horizon line represents a vanishing point from a view point of the ADV based on an initial or default pitch angle of the camera.” *Id.* “One or more lane lines are determined based on the first image via a perception process performed on the first image.” *Id.* “In response to a first input signal (e.g., an up or down arrow key) received from an input device (e.g., keyboard, joystick, voice interactive command), a position of the horizon line is updated based on the first input signal and a position of at least one of the lane lines is updated based on the updated horizon line.” *Id.* Choe consistently describes throughout the specification that the horizon line is determined based on the hardware parameters of the camera. *See, e.g.,* Ex. 1005, ¶¶[0046], [0048], [0055], [0059]. Aside from labeling these as “known elements,” Dr. Darrell fails to provide any evidence of

how a POSITA would expect to replace the hardware-based horizon line determination of Choe with a CNN as described in Workman and achieve a predictable result. These are completely different approaches. In my opinion, there is nothing “predictable” about how the Workman CNN model would be incorporated into Choe’s system.

D. Ground 4

98. Westmacot “relates to image annotation, such as the creation of annotated street road images which can be used (among other things) for training a machine learning road detection component.” Ex. 1010, p.3:5-7.

99. Petitioner basically argues that it would be obvious to use the equations disclosed in Westmacot in Choe’s calibration system because Westmacot uses the equations to compute lane lines. This does not provide any motivation for a designer, starting with Choe, to use the equations provided in Westmacot. Choe does not use the same approach Westmacot uses with respect to detection of lane lines and therefore the equations of Westmacot are not relevant to Choe’s lane line detection. Choe’s system includes a “perception module” with “an object detector to detect an object within the image and a lane detector to detect one or more lanes. One or more lane lines are determined based on the first image by perception module 302.” Ex. 10005, ¶[0048]. Westmacot does not include a perception module because it operates on image sets after they are captured to label the

images for training other models. Ex. 1010, p.3:15-25. In the Westmacot system, “road or lane width may be determined based on one or more manual width adjustment inputs.” *Id.* at 4:24-25. Then, a road normal vector n may be estimated using the normal vector equation of page 5. *See id.* at 5. With this “road structure” defined, border points of the road structure may be determined using the equations the Petition alleges would be obvious to use in Choe. *See* Pet. at 54; Ex. 1010, pp.5-6. However, the Petition does not articulate any reason why these formulas to compute border points in the road structure defined in Westmacot have any applicability to the “lane detector to detect one or more lanes” in Choe’s system. A POSITA would neither be motivated to use these formulas in the system of Choe nor would there be any expectation of success in such a use.

100. The Petitioner also argues that “[s]ince Choe’s camera calibration requires calculating the camera’s rotational orientation (pitch, yaw, and roll), a POSA would have understood that the system would need to compute these camera parameters.” Pet. at 54. But Choe already describes how to obtain pitch, yaw, and roll and does not mention in any way to need to compute camera height and road plane normal the way Westmacot describes it. For example, Choe describes that “[a] user can calibrate any of the sensors such as a pitch angle, a yaw angle, and/or a roll angle using an input device such as a keyboard or joystick to provide incremental calibration feedback, while displaying the calibration result on

a display device concurrently. As a result, the user can visually ‘feel’ the calibration result based on user feedback.” Ex. 10005, ¶[0014]. The approach in Westmacot is simply different and there is no reason or motivation for a POSITA to use the equations of Westmacot in the Choe system.

101. Further, the system described in Westmacot is wholly incompatible with what Choe describes. A POSITA would not be motivated to combine Westmacot with Choe’s camera calibration system because Westmacot expressly notes that its preferred operation works “using only information derived from image data captured by the image capture device ... without the use of any other vehicle sensor and **without any requiring any calibration of the image capture device.**” Ex. 1010, p.3:25-31 (emphasis added). This is precisely the opposite of what Choe teaches. Choe provides a sensor calibration system with a “perception module 302 [that] receives sensor data from a variety of sensors such as camera 211, LIDAR 215, and RADAR 214. ... an image processed by perception module 302 may be utilized by sensor calibration module 308 for sensor calibration.” Ex. 1011 at [0045]-[0046]. Thus, a POSITA would not be motivated to combine Choe’s camera calibration system with Westmacot’s image labeling system that operates “without requiring any calibration of the image capture device.”

102. In addition, Westmacot’s equations do not actually teach “computing at least one camera parameter based on the one or more lines ... wherein the at

least one camera parameter comprises a camera height and road plane normal.”

Ex. 1001 at Claims 1 and 5. The g_i equation shows the computation of road point g_i , not of the camera height h . *See* Ex. 1010, p.6:5. The camera height is an input to that equation. Instead, Westmacot teaches that “[t]he height may be estimated based on an average height of a reconstructed path of the image capture device and a 3D road structure mesh computed from the images.” *Id.* at 7:2-4. This has no application to the system of Choe since Choe does not reconstruct the path of the image capture device. Further, this height is not computed based on any lines identified with a predictive model as claim 5 requires.

I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: November 17, 2025



Scott Andrews