GIBSON, DUNN & CRUTCHER LLP
JOSH KREVITT (SBN 208552)
JKrevitt@gibsondunn.com
STUART M. ROSENBERG (SBN 239926)
SRosenberg@gibsondunn.com
310 University Avenue
Palo Alto, CA 94301
Telephone: 650.849.5300
Facsimile: 650.849.5333

BRIAN A. ROSENTHAL (*pro hac vice*)
BRosenthal@gibsondunn.com
KATHERINE Q. DOMINGUEZ (*pro hac vice*)
KDominguez@gibsondunn.com
ALLEN KATHIR (*pro hac vice*)
AKathir@gibsondunn.com
ALLYSON PARKS (*pro hac vice*)
AParks@gibsondunn.com
200 Park Avenue
New York, NY 10166
Telephone: 212.351.4000
Facsimile: 212.351.4035

GIBSON, DUNN & CRUTCHER LLP
JAYSEN S. CHUNG (SBN 280708)
JSChung@gibsondunn.com
One Embarcadero Center, Suite 2600
San Francisco, CA 94111
Telephone: 415.393.8200
Facsimile: 415.393.8306

NATHANIEL R. SCHARN (SBN 305836)
NScharn@gibsondunn.com
3161 Michelson Drive, Suite 1200
Irvine, CA 92612
Telephone: 949.451.3800
Facsimile: 949.451.4220

HANNAH L. BEDARD (*pro hac vice*)
HBedard@gibsondunn.com
WENDY W. CAI (*pro hac vice*)
WCai@gibsondunn.com
1700 M STREET, NW
Washington, DC 20036
Telephone: 202.955.8500
Facsimile: 202.467.0539

Attorneys for Defendant Samsara Inc.

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

MOTIVE TECHNOLOGIES, INC.,

Plaintiff and Counterclaim Defendant,

v.

SAMSARA INC.,

Defendant and Counterclaim Plaintiff.

CASE NO. 3:24-CV-00902-JD

SAMSARA, INC.'S NOTICE AND MEMORANDUM OF POINTS AND AUTHORITIES IN SUPPORT OF ITS MOTION FOR PARTIAL STAY PENDING *INTER PARTES* REVIEW

Noticed Hearing Date: January 8, 2026
Noticed Hearing Time: 10:00 a.m.
Courtroom 11, 19th Floor
Judge: Honorable James Donato

PLEASE TAKE NOTICE that on November 25, 2025, or as soon thereafter as the parties may be heard before the Honorable Judge James Donato of the United States District Court for the Northern District of California in the San Francisco Courthouse, Courtroom 11, 19th Floor, 450 Golden Gate Avenue, San Francisco, California, Defendant and Counterclaim Plaintiff Samsara, Inc. ("Samsara") will, and hereby does, move this Court for a partial stay of proceedings pending *inter partes* review. This motion is based upon this Notice; the following Memorandum of Points and Authorities; the complete files and records in this action; the argument of counsel; and such other matters as the Court may consider.

Motive Technologies Exhibit 2003, IPR2026-00034
Page 2

TABLE OF CONTENTS

I. INTRODUCTION **1**
II. BACKGROUND **2**
III. LEGAL STANDARD **3**
IV. ARGUMENT **4**
a. A Stay Would Significantly Simplify the Issues 4
i. The IPRs Will Most Likely Result in Invalidation of All Asserted Claims 4
ii. The IPRs Will Narrow the Issues Even if Some Asserted Claims Survive 6
b. The Case Is in Its Early Stages 8
c. A Stay Will Not Unduly Prejudice Motive 8
V. CONCLUSION **10**

TABLE OF AUTHORITIES

Cases

Advanced Analogic Techs., Inc. v. Kinetic Techs., Inc.,
2009 WL 4981164 (N.D. Cal. Dec. 15, 2009) 8

Brightex Bio-Photonics, LLC v. L'Oreal USA, Inc.,
2025 WL 1866339 (N.D. Cal. July 7, 2025) 6

Columbia Ins. Co. v. Simpson Strong-Tie Co.,
703 F. Supp. 3d 1007 (N.D. Cal. 2023) 3, 4, 8, 9, 10

In re Cygnus Telecomms. Tech., LLC Patent Litig.,
385 F. Supp. 2d 1022 (N.D. Cal. 2005) 4

Dialect, LLC v. Google, LLC,
2024 WL 4314206 (N.D. Cal. Sept. 26, 2024) 5

Ethicon, Inc. v. Quigg,
849 F.2d 1422 (Fed. Cir. 1988) 3

Evolutionary Intel., LLC v. Sprint Nextel Corp.,
2014 WL 4802426 (N.D. Cal. Sept. 26, 2014) 7, 8

Finjan, Inc. v. Symantec Corp.,
139 F. Supp. 3d 1032 (N.D. Cal. 2015) 3

LBT IP II LLC v. Uber Techs. Inc.,
2023 WL 322894 (N.D. Cal. Jan. 19, 2023) 7, 8

LELO, Inc. v. Standard Innovation (US) Corp.,
2014 WL 2879851 (N.D. Cal. June 24, 2014) 4, 8, 9

Omega Eng'g, Inc. v. Raytek Corp.,
334 F.3d 1314 (Fed. Cir. 2003) 7

Sonics, Inc. v. Arteris, Inc.,
2013 WL 503091 (N.D. Cal. Feb. 8, 2013) 7

Symantec Corp. v. Zscaler, Inc.,
2018 WL 3539267 (N.D. Cal. July 23, 2018) 9

Topia Tech., Inc. v. Dropbox Inc.,
2023 WL 3437823 (N.D. Cal. May 12, 2023) 5

Statutes

35 U.S.C. §§ 313, 314(b), 316(a)(11) 3

§ 101 of the Patent Act 3

Regulations

37 C.F.R. § 42.107 3

TABLE OF ABBREVIATIONS

ABBREVIATION	DESCRIPTION
'243 patent	U.S. Pat. No. 12,062,243
'276 patent	U.S. Pat. No. 12,136,276
'580 patent	U.S. Pat. No. 11,875,580
IPR	*Inter Partes* Review
Motive	Plaintiff and Counterclaim Defendant Motive Technologies, Inc.
PTAB	Patent Trial and Appeal Board
PTO	U.S. Patent and Trademark Office
Samsara	Defendant and Counterclaim Plaintiff Samsara Inc.

MEMORANDUM OF POINTS AND AUTHORITIES

I. INTRODUCTION

Samsara moves to stay Motive's patent infringement claims pending final resolution of Samsara's IPR petitions challenging all asserted claims.

The PTAB has already instituted IPR proceedings on the sole patent Motive asserted in its original complaint (the '580 patent). Those instituted proceedings alone confirm there is a reasonable likelihood that the claims of the '580 patent are not patentable and will either eliminate all asserted '580 claims or significantly narrow the issues in this case.

In addition, although Motive added two other patents (the '276 and '243 patents) to this case nearly 17 months after Motive filed this lawsuit, Samsara promptly filed IPR petitions on those patents and fully expects the PTAB will institute IPR proceedings on them. The '276 patent is a continuation of the '580 patent, and the claims are so similar that the Patent Office determined claim 1 of each patent to be "not patentably distinct" from each other during prosecution. Consistent with that overlap, Samsara's '276 IPR petition relies on the same primary prior art reference at issue in the already instituted '580 IPR. Institution is also likely for the '243 petition. Motive emphasized the '243 patent's similarity to the '580 and '276 patents when it sought leave to add the '243 patent to this case. *See* Dkt. 96 (Motive's Motion for Leave) at 5. Moreover, the '243 petition relies on previously unconsidered prior art and the PTAB can resolve the IPR well before the August 2027 trial date.

Moreover, Motive cannot rely on the earlier stage of the '276 and '243 IPR petitions to argue that a stay is premature; institution remains pending on those patents only because *Motive* chose to add them to this lawsuit more than 15 months after it was filed, rather than assert them in a new case. Motive should not be permitted to use its late addition of these patents as a reason to avoid a stay. Indeed, to allow otherwise would run directly counter to Motive's prior representation that the timing of its late amendment would not prejudice Samsara or provide Motive a tactical advantage. Absent the addition of these two

patents, Motive would have no arguable basis to oppose a stay in view of the already instituted '580 IPR. *Id.* at 5-6.

Accordingly, while Samsara recognizes that this Court does not typically stay a patent infringement claim before the related IPR is instituted, these facts and the unique procedural posture warrant a stay of all Motive's patent infringement claims now. A stay would streamline this litigation by eliminating Motive's patent infringement claims entirely or narrowing the issues. Further, the bulk of this case has yet to be litigated, discovery is in its early stages, claim construction proceedings have not yet started, and trial will not occur for nearly two years. Finally, Motive would not be unduly prejudiced because the petitions will be resolved before the August 2027 trial date in this case, it was Motive's choice to delay adding the '276 and '243 patents, and monetary damages would be sufficient to remedy any harm to Motive. Thus, with all relevant factors heavily favoring a stay, this motion should be granted.

II. BACKGROUND

On February 15, 2024, Motive filed this action asserting the '580 patent and non-patent causes of action. Dkt. 1. Samsara filed an IPR petition against the '580 patent on February 14, 2025, and the PTAB instituted review on August 27, 2025. IPR2025-00574, Paper 12. The PTAB will render a final written decision within 12 months. *Id.*

On May 15, 2025, Samsara filed an answer and counterclaims to Motive's complaint. Dkt. 94. In its counterclaims, Samsara asserted three patents against Motive. *Id.* To date, Motive has not filed an IPR petition against any of Samsara's patents.

On May 29, 2025, over 15 months after filing this case and only after Samsara filed its counterclaims, Motive moved to amend its complaint to allege infringement of two additional patents: the '243 and '276 patents. *See* Dkt. 96. In seeking leave to add these new patents—over Samsara's opposition—Motive explained the new patents are similar to the '580 patent. *Id.* at 5 ("The new '276 patent is a continuation of the already asserted '580 patent, and all three patents are directed to similar

technologies."). Motive also argued that adding these two patents would not cause Samsara prejudice. Motive's Reply (Dkt. 104) at 4. For example, Motive argued the new patents would "not alter the nature of the case or Samsara's defense strategy." Dkt. 96 at 5 ("Nor can Samsara claim prejudice based on the timing of Motive's motion."). The Court granted Motive's motion for leave to file a Second Amended Complaint, Dkt. 106, and on July 9, 2025, Motive filed the Second Amended Complaint asserting the '243 and '276 patents. Dkt. 107.

Samsara promptly filed IPR petitions on the '276 and '243 patents on October 17 and November 14, respectively—approximately three and four months after the Second Amended Complaint was filed and approximately eight months before the statutory deadline. IPR2026-00034; IPR2026-00108; Exs. A-B. The PTAB will decide whether to institute proceedings within six months and will render a final written decision within 12 months thereafter. 35 U.S.C. §§ 313, 314(b), 316(a)(11); 37 C.F.R. § 42.107. Samsara also moved to dismiss the newly asserted patents as directed toward ineligible subject matter under Section 101 of the Patent Act. Dkt. 112. That motion remains pending.

The case remains in its early stages. Discovery has only just begun—the parties have done no more than exchange initial disclosures, infringement and invalidity contentions, and discovery requests. No depositions have been noticed or taken. The claim construction hearing is set for April 2026, and expert discovery is set to conclude in November 2026. Dkt. 115. Trial is set for August 30, 2027. Dkt. 115.

III. LEGAL STANDARD

"A district court has the inherent power and discretion to stay judicial proceedings pending [IPR]." *Columbia Ins. Co. v. Simpson Strong-Tie Co.*, 703 F. Supp. 3d 1007, 1010 (N.D. Cal. 2023) (citing *Ethicon, Inc. v. Quigg*, 849 F.2d 1422, 1426–27 (Fed. Cir. 1988)); *see also Finjan, Inc. v. Symantec Corp.*, 139 F. Supp. 3d 1032, 1035 (N.D. Cal. 2015) (recognizing "a liberal policy in favor of granting motions to stay proceedings pending the outcome of USPTO . . . proceedings.").

Three factors are relevant in determining whether to grant a stay: "(1) whether discovery is complete and whether a trial date has been set; (2) whether a stay will simplify the issues in question and trial of the case; and (3) whether a stay would unduly prejudice or present a clear tactical disadvantage to the nonmoving party." *Columbia Ins.*, 703 F. Supp. 3d at 1010. "A stay is 'particularly justified where the outcome of the [IPR] would be likely to assist the court in determining patent validity and, if the claims were cancelled in the [IPR], would eliminate the need to try the infringement issue.'" *LELO, Inc. v. Standard Innovation (US) Corp.*, 2014 WL 2879851, at *2 (N.D. Cal. June 24, 2014) (quoting *In re Cygnus Telecomms. Tech., LLC Patent Litig.*, 385 F. Supp. 2d 1022, 1023 (N.D. Cal. 2005)).

IV. ARGUMENT

All three stay factors favor granting this motion. The IPRs will likely eliminate Motive's asserted patents or claims, significantly narrowing the issues for this Court. Moreover, even if Motive's claims were to survive, the record developed in the PTAB will provide the Court with substantive guidance that will streamline claim construction and simplify the issues for trial. In addition, the case is still in its early stages and trial will not occur for nearly two years. Further, a stay would not unduly prejudice or present a tactical disadvantage to Motive.

a. A Stay Would Significantly Simplify the Issues

A stay could eliminate Motive's patent infringement case entirely, and in any event will streamline this case regardless of the IPRs' outcome. Thus, this factor strongly favors a stay. *See LELO*, 2014 WL 2879851, at *3 (noting simplification is the "central purpose" of IPR proceedings).

i. The IPRs Will Most Likely Result in Invalidation of All Asserted Claims

Because Samsara's petitions challenge every asserted patent claim, the IPRs stand to eliminate Motive's patent infringement case entirely and allow the case to proceed solely with Motive's non-patent causes of action and Samsara's patent infringement counterclaims.

IPR proceedings have already been instituted for the '580 patent. The most likely outcome of these instituted proceedings is that all of the '580 patent claims asserted in this case will be held unpatentable. Indeed, the PTAB trial statistics for the most recent data period show that 63% of patents addressed in a Final Written Decision had all claims invalidated. *See* Ex. C (PTAB Trial Statistics).

There is also a high likelihood that the '276 petition will be instituted and materially simplify this case, given the patent's close relationship to the '580 patent and the fact that the same primary prior art reference used in the '580 IPR is also being used to challenge the '276 patent. The '276 and '580 patents share a common specification, and certain independent claims are so similar that the Patent Office deemed them "not patentably distinct from each other." Ex. D, Non-Final Rejection, U.S. Patent Application No. 18/527,266 at 3-4 (USPTO, mailed June 12, 2024). Motive's claim charts further demonstrate the similarity between the claims of these patents, as they cite the same material across the two patents. *Compare* Dkt. 107-2, *with* Dkt. 107-4. Thus, particularly because the same primary prior art is at issue in both petitions, the PTAB's decision to institute the '580 IPR strongly indicates that the '276 IPR will likewise be instituted. *See Topia Tech., Inc. v. Dropbox Inc.,* 2023 WL 3437823, at *4 (N.D. Cal. May 12, 2023) ("[T]he PTAB's institution of review of the '942 patent suggests institution is more likely for the other similar patents asserted in this case.").

The '243 petition is likewise likely to be instituted and to eliminate or materially narrow the issues in this case. Institution is likely because the '243 and '580 patents share overlapping inventors and very similar subject matter. Indeed, Motive cannot deny that the two patents involve similar subject matter, as Motive itself argued they relate to similar technology when it successfully obtained leave to assert the '243 patent in this case. Motive's Motion for Leave (Dkt. 96) at 5; *see Dialect, LLC v. Google, LLC*, 2024 WL 4314206, at *3 (N.D. Cal. Sept. 26, 2024) (noting "the asserted petitions are likely to be instituted" because petitions on "patents with same inventors and similar subject matter" had already been instituted). In addition, Samsara's '243 IPR petition presents the Patent Office with prior art it never considered before

allowing the patent, which further supports granting a pre-institution stay. *See Brightex Bio-Photonics, LLC v. L'Oreal USA, Inc.*, 2025 WL 1866339, at *1 (N.D. Cal. July 7, 2025) (granting pre-institution stay where petition was based on "prior art that was not considered by the PTO during the original examination").

Even though certain changes have taken place at the PTO this year, institution is still likely. PTO Director Squires has made clear that the PTO's decisions over the past year "provid[e] substantial guidance" on how he will "handle discretionary considerations." Ex. E, October 17, 2025 Memorandum from PTO Director Squires at 2. Those decisions illustrate that under the facts here, Samsara's IPR petitions are still likely to be granted. The '276 and '243 patents issued within the past six years and the IPRs would conclude before the scheduled trial date in this case. *See, e.g.*, IPR2025-01023, Paper 12 (finding discretionary denial was not appropriate because IPR would likely conclude before trial and patents issued in 2020 and 2023). Indeed, in the last three weeks, the PTO has instituted 23 IPR petitions involving patents issued within the last six years. *See* IPR2025-00951, IPR2025-00915, IPR2025-00842, IPR2025-00949, IPR2025-00843, IPR2025-00901, IPR2025-00911, IPR2025-00896, IPR2025-00910, IPR2025-00909, IPR2025-00907, IPR2025-00827, IPR2025-00860, IPR2025-00857, IPR2025-00898, IPR2025-00819, IPR2025-00818, IPR2025-00821, IPR2025-00913, IPR2025-01044, IPR2025-01043, IPR2025-00916, IPR2025-00914. These decisions further support that the '276 and '243 IPRs are likely to be instituted and simplify this case. *See Brightex*, 2025 WL 1866339, at *3 (granting stay and noting plaintiff pointed to "nothing specific" indicating that the IPR petition would be subject to discretionary denial).

ii. The IPRs Will Narrow the Issues Even if Some Asserted Claims Survive

Even if any asserted claims ultimately survive, the IPRs will still narrow the issues. Given the closeness of the prior art, in order to overcome it, Motive would need to make narrowing admissions that would bind Motive in this litigation and limit the scope of the claims. Moreover, any statements made by

Motive during the IPRs will become part of the patents' prosecution histories and will aid the Court during claim construction. *See, e.g.*, *Evolutionary Intel., LLC v. Sprint Nextel Corp*., 2014 WL 4802426, at *4 (N.D. Cal. Sept. 26, 2014) ("Statements made by [patentee] during the IPR could disclaim claim scope, aid the court in understanding the meaning of the terms, or otherwise affect the interpretation of key terms."). Thus, staying this case while allowing the IPRs to proceed will develop the intrinsic record, and, to the extent any claims survive, reduce disputes over the meaning and scope of key terms, streamline Markman proceedings, and focus fact and expert discovery. *Id.*

Additionally, a stay would further streamline the issues because Samsara has stipulated to be estopped from asserting invalidity arguments it raised or could have reasonably raised before the PTO for the '580 patent. *See LELO*, 2013 WL 503091, at *3 (noting estoppel "heavily tips the scale in favor of granting the stay"); *LBT IP II LLC v. Uber Techs. Inc.*, 2023 WL 322894, at *2 (N.D. Cal. Jan. 19, 2023) ("Even if all of the asserted claims survive, [the defendant] will be bound by the estoppel provisions, which would also simplify the issues before me.").

Even in the unlikely event that the '276 IPR or '243 IPR is not instituted, the PTAB's final written decision on the '580 patent would still significantly simplify this case. The '580 patent is related to the '276 patent and involves similar technology as compared to the '243 patent. *See Sonics, Inc. v. Arteris, Inc.*, 2013 WL 503091, at *3 (N.D. Cal. Feb. 8, 2013) (staying case as to all patents due to "overlapping issues between the reexamined patents and other non-reexamined patents-in-suit"). As a result, it is highly likely that arguments pertaining to the '580 patent will impact the '276 patent, because the prosecution history of a related patent is relevant to claim construction. *Omega Eng'g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1334 (Fed. Cir. 2003) ("The patentee made a clear and unmistakable disclaimer of claim scope in its prosecution of the parent '880 patent, and we presume, unless otherwise compelled, that the same claim term in the same patent or related patents carries the same construed meaning."). Given the ability of an instituted patent's petition to impact pre-institution patents, courts in this District have stayed a litigation

as to all patent claims even where IPR proceedings have been instituted only on some of the patents. *See, e.g.*, *Columbia Ins.*, 703 F. Supp. 3d at 1012 (staying case prior to institution decision where related patent had already been partially invalidated); *Evolutionary Intel.*, 2014 WL 4802426, at *4 (collecting cases). In sum, the case-simplification factor heavily favors a stay.

b. The Case Is in Its Early Stages

The early posture of this case strongly favors a stay. Discovery has only recently begun—the parties have exchanged only initial disclosures, infringement and invalidity contentions, and discovery requests. The parties have not yet started claim construction. No depositions have been taken, and expert discovery has not started. Consistent with these circumstances, courts have found that the early posture of this case favors a stay. *See LELO*, 2014 WL 2879851, at *3 ("While not a newborn, this case is still in infancy."); *see also LBT*, 2023 WL 322894, at *1 (finding early stage of case favored stay even where opening claim construction brief had already been filed).

Moreover, trial is not scheduled to occur for another 21 months, and courts in this District have found the timing favored a stay even in cases where trial was much more imminent. *See LELO*, 2014 WL 2879851, at *3 (13 months); *Advanced Analogic Techs., Inc. v. Kinetic Techs., Inc.*, 2009 WL 4981164, at *2 (N.D. Cal. Dec. 15, 2009) (15 months). At bottom, there is substantial work that remains to be done before this case proceeds to trial, including the most resource and time intensive parts of a case (*e.g.*, claim construction, fact discovery, depositions, expert discovery, dispositive motions, and trial). The status of the case therefore strongly favors a stay.

c. A Stay Will Not Unduly Prejudice Motive

Staying Motive's patent infringement claims pending the outcome of the IPR proceedings will not unduly prejudice Motive. Courts typically consider four factors in evaluating prejudice: "(1) the timing of the petition for review; (2) the timing of the request for the stay; (3) the status of review proceedings;

and (4) the relationship of the parties." *Symantec Corp. v. Zscaler, Inc.*, 2018 WL 3539267, at *3 (N.D. Cal. July 23, 2018). Here, all four factors strongly favor a stay.

Both the timing of the petition and the timing of the request factors support a stay. Samsara filed the '276 and '243 petitions promptly after Motive added those patents—well ahead of the one-year statutory deadline—and brings this motion within seven business days of filing the last of the petitions. Samsara also filed these petitions and brought this request for a stay during the early stages of this case. *Supra* 6-7. To the extent Motive argues that there was any delay, that delay was caused by Motive's late addition of the '276 and '243 patents in this lawsuit. Both patents issued by November 5, 2024, yet Motive chose not to immediately assert them in a new lawsuit. Instead, it waited until the stay in this case was lifted (in March), and then, on May 29, 2025—more than 15 months after filing this lawsuit—it moved to amend the complaint to add both patents. Dkt. 96. After Motive filed its Second Amended Complaint on July 9, 2025 (Dkt. 107), Samsara promptly filed IPR petitions.

The status of the IPRs also favors a stay. *Supra* 4-5. Samsara's '580 IPR has been instituted and the other two IPRs are highly likely to be instituted and completed before the August 2027 trial date. *Id.* Further, the lack of institution decisions on the '276 and '243 petitions is entirely due to Motive's belated assertion of the '243 and '276 patents. Allowing Motive to benefit from its staggered assertion of patents would unduly prejudice Samsara—the opposite of what Motive argued would occur if the Court permitted it to add these two patents to this case. Thus, although this Court typically waits until after institution to grant a stay, the unique circumstances of this case warrant a stay now.

Lastly, although Motive and Samsara are competitors, this relationship "is not on its own enough to outweigh the other factors favoring the granting of the requested stay." *Columbia Ins.*, 703 F. Supp. 3d at 1013; *see also LELO*, 2014 WL 2879851, at *4 ("[W]hile some competition appears to exist between the parties, the overall state of affairs does not raise the specter of undue prejudice or clear disadvantage to [the plaintiff]."). Here, Motive has not sought a preliminary injunction and monetary damages would

be adequate to remedy any harm, should Motive ultimately succeed on the patent claims. *See Columbia Ins.*, 703 F. Supp. 3d at 1013. Likewise, any argument that a stay would unduly prejudice Motive is undermined by Motive's delay in asserting the '276 and '243 patents.

Thus, the lack of undue prejudice to Motive strongly favors a stay.

V. CONCLUSION

For the foregoing reasons, this Court should stay Motive's patent infringement claims pending final resolution of Samsara's IPR petitions concerning Motive's three asserted patents, including any related appeals.

DATED: November 25, 2025

Respectfully submitted,

/s/ Josh A. Krevitt

Josh A. Krevitt (SBN 208552)
Brian A. Rosenthal (*pro hac vice*)
Katherine Q. Dominguez (*pro hac vice*)
Allen Kathir (*pro hac vice*)
Allyson Parks (*pro hac vice*)
GIBSON, DUNN & CRUTCHER LLP
200 Park Avenue
New York, NY 10166-0193
Telephone: 212.351.3961
Facsimile: 212.716.0781
jkrevitt@gibsondunn.com
brosenthal@gibsondunn.com
kdominguez@gibsondunn.com
akathir@gibsondunn.com
aparks@gibsondunn.com

Jaysen S. Chung (SBN 280708)
GIBSON, DUNN & CRUTCHER LLP
One Embarcadero Center, Suite 2600
San Francisco, CA 94111-3715
Telephone: 415.393.8200
Facsimile: 415.393.8306
jschung@gibsondunn.com

Stuart M. Rosenberg (SBN 239926)
srosenberg@gibsondunn.com
GIBSON, DUNN & CRUTCHER LLP
310 University Avenue
Palo Alto, CA 94301
Telephone: 650.849.5300
Facsimile: 650.849.5333

Nathaniel R. Scharn (SBN 305836)
GIBSON, DUNN & CRUTCHER LLP
3161 Michelson Drive, Suite 1200
Irvine, CA 92612-4412
Telephone: 949.451.3800
Facsimile: 949.451.4220
nscharn@gibsondunn.com

Hannah L. Bedard (*pro hac vice*)
Wendy Cai (*pro hac vice*)
GIBSON, DUNN & CRUTCHER LLP
1700 M Street, N.W.
Washington, DC 20036-4504
Telephone: 202.955.8500
Facsimile: 202.467.0539
hbedard@gibsondunn.com
wcai@gibsondunn.com

CERTIFICATE OF SERVICE

I hereby certify that counsel of record who are deemed to have consented to electronic service are being served on November 25, 2025, with a copy of this document via electronic mail.

/s/ Josh A. Krevitt
Josh A. Krevitt

GIBSON, DUNN & CRUTCHER LLP
JOSH KREVITT (SBN 208552)
JKrevitt@gibsondunn.com
STUART M. ROSENBERG (SBN 239926)
SRosenberg@gibsondunn.com
310 University Avenue
Palo Alto, CA 94301
Telephone: 650.849.5300
Facsimile: 650.849.5333

BRIAN A. ROSENTHAL (*pro hac vice*)
BRosenthal@gibsondunn.com
KATHERINE Q. DOMINGUEZ (*pro hac vice*)
KDominguez@gibsondunn.com
ALLEN KATHIR (*pro hac vice*)
AKathir@gibsondunn.com
ALLYSON PARKS (*pro hac vice*)
AParks@gibsondunn.com
200 Park Avenue
New York, NY 10166
Telephone: 212.351.4000
Facsimile: 212.351.4035

GIBSON, DUNN & CRUTCHER LLP
JAYSEN S. CHUNG (SBN 280708)
JSChung@gibsondunn.com
One Embarcadero Center, Suite 2600
San Francisco, CA 94111
Telephone: 415.393.8200
Facsimile: 415.393.8306

NATHANIEL R. SCHARN (SBN 305836)
NScharn@gibsondunn.com
3161 Michelson Drive, Suite 1200
Irvine, CA 92612
Telephone: 949.451.3800
Facsimile: 949.451.4220

HANNAH L. BEDARD (*pro hac vice*)
HBedard@gibsondunn.com
WENDY W. CAI (*pro hac vice*)
WCai@gibsondunn.com
1700 M STREET, NW
Washington, DC 20036
Telephone: 202.955.8500
Facsimile: 202.467.0539

Attorneys for Defendant Samsara Inc.

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

MOTIVE TECHNOLOGIES, INC.,

Plaintiff and Counterclaim Defendant,

v.

SAMSARA INC.,

Defendant and Counterclaim Plaintiff.

CASE NO. 3:24-CV-00902-JD

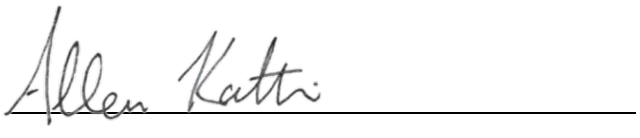
DECLARATION OF ALLEN KATHIR IN SUPPORT OF SAMSARA, INC.'S MOTION TO STAY MOTIVE'S PATENT INFRINGEMENT CLAIMS PENDING *INTER PARTES* REVIEW

I, Allen Kathir, the undersigned, pursuant to 28 U.S.C. § 1746, declare as follows:

1. I am an associate at the law firm of Gibson Dunn & Crutcher LLP, counsel of record for and Counterclaim Plaintiff Samsara Inc. ("Samsara"). I am admitted and in good standing to practice law in the state of New York. I have personal knowledge of the matters set forth below and, if called and sworn as a witness, I could and would testify competently to the facts set forth herein.
2. I respectfully submit this declaration pursuant to Civil L.R. 7-5 in support of Samsara's Motive to Stay Motive's Patent Infringement Claims pending *Inter Partes* Review.
3. Attached hereto as Exhibit A is a true and correct copy of the *inter partes* review petition in proceeding number IPR2026-00034 regarding U.S. patent No. 12,136,276.
4. Attached hereto as Exhibit B is a true and correct copy of the *inter partes* review petition in proceeding number IPR2026-00108 regarding U.S. patent No. 12,062,243.
5. Attached hereto as Exhibit C is a true and correct copy of inter partes review statistics from Oct. 1, 2024 to Sept. 30, 2025, obtained from https://www.uspto.gov/sites/default/files/documents/Trial_StatsFY25_Q4.pdf.
6. Attached hereto as Exhibit D is a true and correct copy of a June 12, 2024 office action from the prosecution history of U.S. patent No. 12,136,276.
7. Attached hereto as Exhibit E is a true and correct copy of an October 17, 2025 memorandum from Patent and Trademark Office Director John A. Squires.
8. I declare under penalty of perjury under the laws of the United States that the foregoing is true and correct.

Executed in New York, New York.

DATED: November 25, 2025



Allen Kathir

EXHIBIT A

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSARA INC.

Petitioner

v.

MOTIVE TECHNOLOGIES, INC.

Patent Owner

Case IPR2026-00034
U.S. Patent No. 12,136,276

PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 12,136,276

Mail Stop "PATENT BOARD"
Patent Trial and Appeal Board
U.S. Patent & Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

I. INTRODUCTION 1
II. U.S. PATENT NO. 12,136,276 2
A. Overview 2
B. Prosecution History 3
C. Level of Ordinary Skill in the Art 3
D. Claim Construction 4
III. IDENTIFICATION OF THE CHALLENGE (37 C.F.R. § 42.104(B)) 4
A. Prior Art 4
B. Grounds 5
IV. OVERVIEW OF THE PRIOR ART 5
A. Choe 5
B. Westmacot 8
C. Davies 11
D. Tal 13
E. Kuehnle 14
V. GROUND 1: CHOE RENDERS OBVIOUS CLAIMS 1-3, 5-10, 12-16 AND 18-20. 14
A. Independent Claim 1 15
1. [1Pre] "A method comprising" 15
2. [1A] "receiving an image of a roadway recorded by a camera device installed within a vehicle" 15
3. [1B] "detecting a horizon line in the image" 16
4. [1C] "overlaying a line on the image to generate an overlaid image" 19
5. [1D] "transmitting the overlaid image to a computing device over a network" 20
6. [1E] receiving a modification of the line from the computing device, the modification comprising a new line at a second position; 23

7. [1F] computing a camera parameter based on the new line 25

8. [1G] transmitting data representing the camera parameter to the camera device 27

B. Claim 2: "the camera parameter comprises one of camera height, viewing angle, and road plane normal." 28

C. Claim 3: "receiving the image of the roadway comprises receiving an image of the roadway while the vehicle is moving." 29

D. Claim 5: "transmitting the overlaid image to the computing device over the network comprises transmitting the overlaid image to a web-based application." 29

E. Claim 6: "receiving the modification of the line from the computing device comprises receiving an identification of a new horizon line by a user of the computing device." 31

F. Claim 7: "transmitting the data representing the camera parameter to the camera device further comprises recomputing the camera parameter based on the modification of the line." 33

G. Claims 8-10, 12-16 18-20 are obvious over Choe. 35

VI. GROUND 2: CHOE IN VIEW OF DAVIES RENDERS OBVIOUS CLAIMS 1-3, 5-10, 12-16 AND 18-20. 36

A. A POSA would have been motivated to combine Choe with Davies. 36

1. Choe 36

2. Davies 37

3. *KSR* 38

B. Claim 1 41

1. 1[Pre]-1[C], 1[F], 1[G] 41

2. 1[D]: "transmitting the overlaid image to a computing device over a network" 42

a. Choe 42

b. Davies 42

c. *KSR* 43

3. 1[E]: "receiving a modification of the line … comprising a new line at a second position" 44
a. Choe 44
b. Davies 45
c. *KSR* 45
C. Claims 2-3 and 5-7 46
1. Claim 2 (camera height, viewing angle, road-plane normal). 46
2. Claim 3 (image received while the vehicle is moving). 47
3. Claim 5 (transmit to a web-based application). 47
a. Choe 47
b. Davies 48
c. *KSR* 48
4. Claim 6 (receiving a user identification of a new horizon line). 49
a. Choe 49
b. Davies 50
c. KSR 50
5. Claim 7 (recomputing the camera parameter based on the modification). 51
a. Choe 51
b. Davies 51
c. *KSR* 51
D. Claims 8-10, 12-16, and 18-20 52
VII. GROUNDS 3A/3B: CHOE IN VIEW OF KUEHNLE (GROUND 3A) AND CHOE IN VIEW OF DAVIES AND FURTHER IN VIEW OF KUEHNLE (GROUND 3B) RENDER OBVIOUS CLAIMS 4, 11, AND 17 53
A. Claims 4, 11, and 17: "receiving the image of the roadway while the vehicle is moving comprises detecting that the vehicle is traveling above a pre-defined speed or for a pre-defined duration." (Grounds 3A/3B) 54

1. Choe (Ground 3A) and Choe in view of Davies (Ground 3B) 54
2. Kuehnle 55
3. *KSR* 56

VIII. GROUND 4: WESTMACOT IN VIEW OF TAL RENDERS OBVIOUS CLAIMS 1-20. 57

A. A POSA would have been motivated to combine Westmacot and Tal. 57
1. Westmacot 57
2. Tal 58
3. *KSR* 58

B. Independent Claim 1 60
1. [1Pre] "A method comprising" 60
2. [1A] "receiving an image of a roadway recorded by a camera device installed within a vehicle" 60
3. [1B] "detecting a horizon line in the image" 61
4. [1C] "overlaying a line on the image to generate an overlaid image" 62
5. [1D] "transmitting the overlaid image to a computing device over a network" 65
   a. Westmacot 65
   b. Tal 65
   c. *KSR* 66
6. [1E] "receiving a modification of the line from the computing device, the modification comprising a new line at a second position" 67
   a. Westmacot 67
   b. Tal 68
   c. *KSR* 68
7. [1F] computing a camera parameter based on the new line 69
   a. Westmacot 69

b. Tal 70
c. *KSR* 70
8. [1G] transmitting data representing the camera parameter to the camera device 71
a. Westmacot 71
b. Tal 72
c. *KSR* 72
C. Claim 2 73
D. Claim 3 74
E. Claim 4 74
F. Claim 5 75
G. Claim 6 77
H. Claim 7 78
I. Claims 8-20 are obvious over Westmacot and Tal. 79
IX. NO OBJECTIVE INDICIA OF PATENTABILITY 79
X. MANDATORY NOTICES 80
A. Real Party in Interest 80
B. Related Matters 80
A. Lead and Backup Counsel 80
XI. GROUNDS FOR STANDING 81
XII. CONCLUSION 82

TABLE OF AUTHORITIES

Page(s)

Cases

ACCO Brands Corp. v. Fellowes, Inc.,
813 F.3d 1361 (Fed. Cir. 2016) 24

In re Applied Materials, Inc.,
692 F.3d 1289 (Fed. Cir. 2012) 38

Boston Scientific Scimed, Inc. v. Cordis Corp.,
554 F.3d 982 (Fed. Cir. 2009) 32, 38

CRFD Rsch., Inc. v. Matal,
876 F.3d 1330 (Fed. Cir. 2017) 24

DyStar Textilfarben GmbH v. C.H. Patrick Co.,
464 F.3d 1356 (Fed. Cir. 2006) 40

Leapfrog Enters., Inc. v. Fisher-Price, Inc.,
485 F.3d 1157 (Fed. Cir. 2007) 31

Perfect Web Techs., Inc. v. InfoUSA, Inc.,
587 F.3d 1324 (Fed. Cir. 2009) 24, 31

Phillips v. AWH Corp.,
415 F.3d 1303 (Fed. Cir. 2005) 4

Realtime Data, LLC v. Iancu,
No. 18-1154, slip op. (Fed. Cir. Jan. 10, 2019) 32

Uber Techs., Inc. v. X One, Inc.,
957 F.3d 1334 (Fed. Cir. 2020) 24, 33

Wyers v. Master Lock Co.,
616 F.3d 1231 (Fed. Cir. 2010) .....*passim*

Other Authorities

37 C.F.R. § 42.104(B).....4

U.S. Patent No. 11,875,580.....3, 4, 80

U.S. Patent No. 12,062,243.....80

U.S. Patent No. 12,136,276.....*passim*

V.A.1-4, 7, 8.....41, 52

PETITIONER'S EXHIBIT LIST

Exhibit	Description
1001	U.S. Patent No. 12,136,276 (Hassan)
1002	Prosecution History of U.S. Patent No. 12,136,276
1003	Declaration of Dr. Trevor Darrell
1004	*Curriculum Vitae* of Dr. Trevor Darrell
1005	U.S. Publication No. 2020/0410704 (Choe)
1006	Prosecution History Figures of U.S. Publication No. 2020/0410704 to Choe
1007	U.S. Publication No. 2014/0240500 (Davies)
1008	International Publication No. WO2019/175286 (Westmacot)
1009	U.S. Publication No. 2022/0019829 (Tal)
1010	International Publication No. WO2009/027090 (Kuehnle)
1011	U.S. Publication No. 2019/0034740 (Kwant1)
1012	U.S. Publication No. 2019/0102674 (Kwant2)
1013	*Motive Techs., Inc. v. Samsara Inc.*, No. 24-cv-00902, Complaint (NDCA July 9, 2025)
1014	Russell, B.C. & Torralba, A., "Building a database of 3D scenes from user annotations," *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2711-2718 (2009)
1015	Bartl, V. & Herout, A., "Fully Automatic Horizon Estimation for Surveillance Cameras," *2017 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pp. 1-8 (2017)
1016	U.S. Patent No. 9,201,421 to Fairfield et al.
1017	Chinese Publication No. 112509054 to Chongqing, including certified English-language translation
1018	Geiger, A., et al., "Vision Meets Robotics: The KITTI Dataset," *Int'l J. Robotics Res.*, 32(11):1231-1237 (Sept. 2013)

Exhibit	Description
1019	Cordts, M., et al., "The Cityscapes Dataset for Semantic Urban Scene Understanding," *2016 IEEE Conf. on Computer Vision & Pattern Recognition (CVPR)*, 3213-3223 (2016)
1020	Kato, S., et al., "Autoware on Board: Enabling Autonomous Vehicles with Embedded Systems," *Proc. 2018 ACM/IEEE Int'l Conf. on Cyber-Physical Systems*, pp. 287-296 (April 2018)
1021	U.S. Patent No. 8,769,396 to Chen et al.
1022	Graf, G., et al., "The Predictive Corridor: A Virtual Augmented Driving Assistance System for Teleoperated Autonomous Vehicles," International Conference on Artificial Reality and Telexistence Eurographics Symposium on Virtual Environments, pp. 61-69 (2020)
1023	U.S. Patent No. 5,652,849 to Conway et al.
1024	U.S. Patent No. 9,775,682 to Quaid et al.
1025	U.S. Patent No. 9,916,703 to Levinson et al.
1026	Waymo Safety Report: On the Road to Fully Self-Driving (Feb. 2021)
1027	Huang, X., et al., "The ApolloScape Open Dataset for Autonomous Driving and Its Application to Network Benchmarking," *Proc. IEEE CVPR Workshops* (2018)
1028	U.S. Patent No. 9,905,949 to Hartmann
1029	Saparia, S., et al., "Active Safety System for Semi-Autonomous Teleoperated Vehicles," *IEEE* (2021)
1030	U.S. Patent No. 10,027,031 to Arai et al.
1031	Liu, S., et al., "Edge Computing for Autonomous Driving: Opportunities and Challenges," *IEEE*, (99):1-20 (2019)
1032	U.S. Patent No. 9,792,569 to Ikawa
1033	Lee, J., et al., "Online Extrinsic Camera Calibration Using Lane Boundary Observations," pp. 1-6 (2020)
1034	Lee, J., et al., "CTRL-C: Camera Calibration Transformer With Line-Classification," *Proceedings of the IEEE/CVF Int'l Conf. on Computer Vision*, pp. 1-14 (2021)

Exhibit	Description
1035	Andrés M. Lopez et al., *Deep Single Image Camera Calibration With Radial Distortion*, in *Proceedings of the IEEE/CVF Conf. on Computer Vision & Pattern Recognition* 9290 (2019)
1036	Alex Davies, *Nissan's Path to Self-Driving Cars? Humans in Call Centers*, WIRED (Jan. 5, 2017), https://www.wired.com/2017/01/nissans-self-driving-teleoperation/
1037	Aarian Marshall, *Self-Driving Cars Have a Secret Weapon: Remote Control*, WIRED (Feb. 1, 2018), https://www.wired.com/story/phantom-teleops/
1038	Andreas Geiger, Philip Lenz & Raquel Urtasun, *Are We Ready for Autonomous Driving? The KITTI Vision Benchmark Suite*, in 2012 IEEE Conf. on Computer Vision & Pattern Recognition (CVPR) 3354 (2012), http://www.cvlibs.net/publications/Geiger2012CVPR.pdf
1039	*The KITTI Vision Benchmark Suite—Road/Lane Detection Evaluation*, CVlibs (dataset site), http://www.cvlibs.net/datasets/kitti/eval_road.php
1040	*The KITTI Vision Benchmark Suite* (home page), CVlibs, https://www.cvlibs.net/datasets/kitti/
1041	*The Cityscapes Dataset*, **Cityscapes** (dataset site), https://www.cityscapes-dataset.com/
1042	U.S. Dep't of Transp., Nat'l Highway Traffic Safety Admin., Assessment of Safety Standards for Automotive Electronic Control Systems (2016), https://www.nhtsa.gov/sites/nhtsa.gov/files/812285_electronicsreliabilityreport.pdf
1043	J. Brewer et al., *Functional Safety Assessment of a Generic Automated Lane Centering System and Related Foundational Vehicle Systems* (Volpe Nat'l Transp. Sys. Ctr. & NHTSA, Rep. No. DOT HS 812 572, Aug. 2018), https://www.nhtsa.gov/sites/nhtsa.gov/files/documents/13496_812572_alcsynthesis_080318.pdf
1044	Nat'l Highway Traffic Safety Admin., Human Factors Evaluation of Level 2 and Level 3 Automated Driving Concepts (DOT HS 812 182, Aug. 2015),

Exhibit	Description
	https://www.nhtsa.gov/sites/nhtsa.gov/files/812182_humanfactorseval-l2l3-automdrivingconcepts.pdf
1045	Nat'l Highway Traffic Safety Admin., Human Factors Design Guidance for Level 2 and Level 3 Automated Driving Concepts (DOT HS 812 555, Aug. 2018), https://www.nhtsa.gov/sites/nhtsa.gov/files/documents/13494_812555_l2l3automationhfguidance.pdf
1046	U.S. Dep't of Transp. & Nat'l Highway Traffic Safety Admin., Federal Automated Vehicles Policy: Accelerating the Next Revolution in Roadway Safety (Sept. 2016), https://www.transportation.gov/AV/federal-automated-vehicles-policy-september-2016
1047	Aarian Marshall, *Sacramento Eases Into the Self-Driving Scene*, WIRED (Aug. 1, 2018), https://www.wired.com/story/sacramento-phantom-auto-self-driving-car-partnership/
1048	Apollo, Dreamview Plus: README (2020), https://apollo.baidu.com/docs/apollo/9.x/md_modules_2dreamview__plus_2README.html
1049	Cruise, Webviz (2019), https://webviz.io; https://web.archive.org/web/20191229152903/https://webviz.io/
1050	Cruise, Webviz GitHub Repository (2019), https://github.com/cruise-automation/webviz
1051	Kirsten Korosec, *Cruise Is Sharing Its Data Visualization Tool with Robotics Geeks Everywhere*, TechCrunch (June 18, 2019), https://techcrunch.com/2019/06/18/cruise-is-sharing-its-data-visualization-tool-with-robotics-geeks-everywhere
1052	Robot Web Tools, https://robotwebtools.github.io
1053	Amazon Web Servs., Inc., Amazon QuickSight—Fast, Easy Business Analytics for the Cloud, https://aws.amazon.com/quicksight; https://web.archive.org/web/20151118211525/https://aws.amazon.com/quicksight/
1054	U.S. Patent No. 7,650,210 to Breed
1055	U.S. Publication No. 2013/0246135 to Wang

Exhibit	Description
1056	U.S. Patent Application Publication No. 2017/0010106 to Shashua
1057	U.S. Patent No. 9,665,100 to Shashua
1058	U.S. Patent Application Publication No. 2016/0214533 to Doyle
1059	U.S. Patent No. 8,612,136 to Levine
1060	Emami, Yousef, et al. "Human-in-the-loop machine learning for safe and ethical autonomous vehicles: Principles, challenges, and opportunities." *arXiv preprint arXiv:2408.12548* (2024).
1061	Chen, Qi, et al. "F-cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3D point clouds." *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*. 2019.
1062	Xu, Jiaxuan, et al. "An automated learning-based procedure for large-scale vehicle dynamics modeling on baidu apollo platform." *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019.
1063	Jiang, Shu, et al. "DRF: A framework for high-accuracy autonomous driving vehicle modeling." *arXiv preprint arXiv:2011.00646* (2020).
1064	Tang, Jie, et al. "LoPECS: A low-power edge computing system for real-time autonomous driving services." *IEEE Access* 8 (2020): 30467-30479.

I. INTRODUCTION

U.S. Patent No. 12,136,276 ("the '276 patent") claims nothing more than routine camera-calibration techniques long used in autonomous systems/ADAS (advanced driver assist systems).

For Ground 1, Choe renders claims 1–20 obvious by teaching the complete loop: roadway image capture; horizon/lane detection and overlay; user edits; recomputation of pitch/yaw/roll; and client–server calibration with results uploaded to the vehicle.

Under Ground 2, any alleged transmission or remote processing gap in Choe is closed by Davies's conventional, bidirectional transmission of high-resolution/annotated imagery to remote GUIs with processed outputs returned to update local parameters.

Building on those showings, Grounds 3A and 3B render claims 4, 11, and 17 obvious: Kuehnle's motion-based acceptance (using frames only above a predefined speed/duration) applied to Choe's (and Choe+Davies) in-drive capture yields predictable gains: better accuracy, less drift, and more robust lane detection.

Independently, for Ground 4, Westmacot's line-editing/parameter-recalculation in view of Tal's networked pipeline for raw/processed (including overlaid) images sent to a server and disseminated back to clients yields the same result.

The references collectively disclose each element of the challenged claims, and provide clear reasons to combine: improving calibration accuracy, reducing sensor drift, and enhancing robustness of lane detection. For these reasons, the Board should institute trial and find the challenged claims of the '276 patent unpatentable.

II. U.S. PATENT NO. 12,136,276

A. Overview

The '276 patent, titled "Camera Initialization for Lane Detection and Distance Estimation using Single-View Geometry," describes conventional lane-detection/calibration in automotive machine learning systems via a "network-based initialization system" that ingests video, identifies lane/horizon lines, and computes camera parameters. EX1001, 1:14–16, 1:43–51; EX1003, ¶28. The '276 patent uses four blocks: device 230 (vehicle camera/processor), camera-initialization service 232, annotator device 234, and initialization service 236. *Id*., 5:1–7:18. Device 230 sends frames, service 232 detects lines, computes parameters (height, viewing angle, road-plane normal), overlays them, then forwards the overlaid images to annotator device 234, a workstation/web app, for confirmation or edits. *Id*., 5:1–59, 6:6–16. If edited, annotator 234 may send a revised image to service 236 (deep learning/artificial intelligence) for parameter prediction and return updated parameters via service 232 to device 230. *Id.*, 6:33–54, 7:9–18.

B. Prosecution History

The prosecution of the '276 patent involved one Office Action and response, with no amendments. The Examiner rejected claims 21–22 and 28 for non-statutory double patenting over U.S. Patent No. 11,875,580 ("the '580 patent") and claims 21–40 under §103 over Kwant1 (US 2019/0034740) and Kwant2 (US 2019/0102674). EX1002, 119, 122; EX1003, ¶¶93-95.

In response, the Applicant filed a terminal disclaimer and argued: (1) Kwant1 lacks visualization or overlay of a horizon line, showing only object annotations (EX1002, 137); (2) Kwant2 teaches model training, not modification or verification of a line (EX1002, 138–139); and (3) Kwant2 computes only neural-network weights, not camera parameters (EX1002, 139). EX1003, ¶¶93-95.

The Examiner then allowed all claims. As this Petition shows, Choe, Davies, Westmacot, and Tal disclose the very features leading to allowance. EX1003, ¶¶430-434. None were applied during prosecution, and only Westmacot was cited in the IDS.

C. Level of Ordinary Skill in the Art

A person of ordinary skill in the art ("POSA") at the time of the alleged invention would have had a Bachelor's in Computer Science with a focus on computer/machine vision, and at least two years of experience with each of: (1) machine learning (including neural network) methods as applied to

machine/computer vision and (2) classical machine/computer vision algorithms (e.g., edge detection, line identification, computation of camera parameters from horizon line, etc.). EX1003, ¶98. Additional education could serve as a substitute for the experience requirement. *Id.* This petition is supported by an expert declaration of Dr. Trevor Darrell. See EX1003, ¶¶1-27.

D. Claim Construction

Samsara has construed the terms according to their ordinary and customary meanings as understood by a POSA in view of the specification. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312-13 (Fed. Cir. 2005).

III. IDENTIFICATION OF THE CHALLENGE (37 C.F.R. § 42.104(B))

A. Prior Art

The earliest priority date the '276 patent can claim is October 4, 2021—the filing date of its parent application. The Grounds in this Petition rely on the following prior-art references:

Reference	Exhibit(s)	Date	Prior Art
Choe (and its Patent Center filed drawings presented merely for enhanced clarity)	EX1005 EX1006	Published Dec. 31, 2020	§102(a)(1)-(2)
Westmacot	EX1008	Published Sept. 19, 2019	§102(a)(1)-(2)
Davies	EX1007	Published August 28, 2014	§102(a)(1)-(2)

Reference	Exhibit(s)	Date	Prior Art
Tal	EX1009	Filed Jul 15, 2020	§102(a)(2)
Kuehnle	EX1010	Published March 5, 2009	§102(a)(1)-(2)

B. Grounds

Ground	Basis	Reference(s)	Claims
1	§103	Choe	1-3, 5-10, 12-16, 18-20
2	§103	Choe and Davies	1-3, 5-10, 12-16, 18-20
3A/3B	§103	3A: Choe in view Kuehnle; 3B: Choe in view of Davies and further in view of Kuehnle	4, 11, 17
4	§103	Westmacot and Tal	1-20

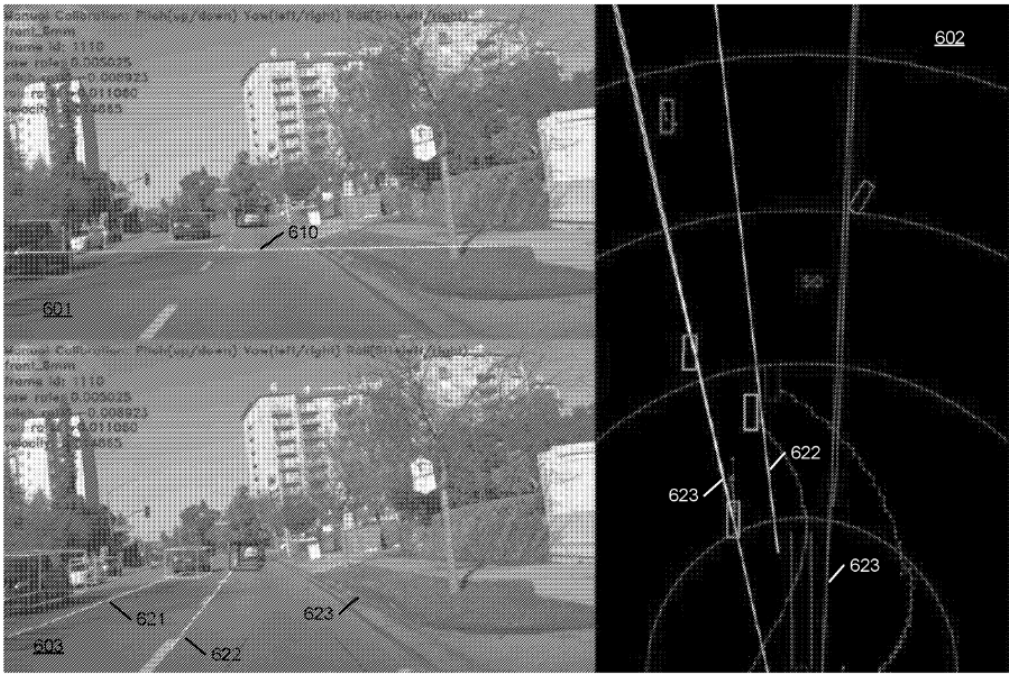
IV. OVERVIEW OF THE PRIOR ART

As the following prior art establishes, manually calibrating a camera to detect traffic lanes was well known. EX1003, ¶29. As part of the calibration process, it was known how to overlay video frames with a horizon line, and then use that horizon line or an adjusted horizon line to calculate the camera's pitch, yaw, and roll. EX1003, ¶¶29-53.

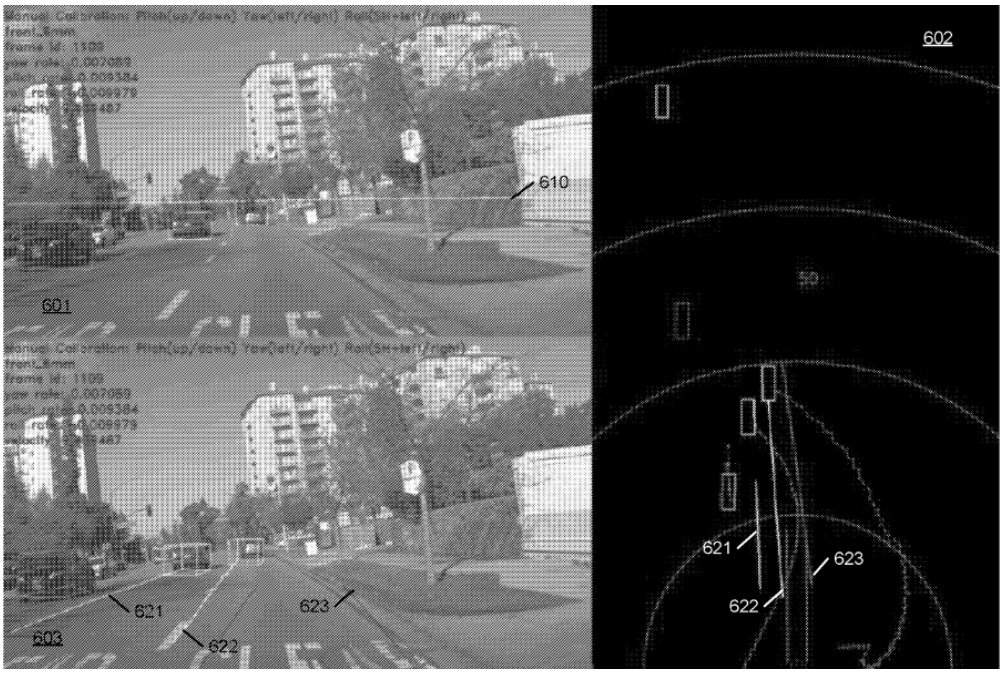
A. Choe

Choe is directed to "calibrating a sensor system of an autonomous driving vehicle," noting sensors require periodic calibration but "a lack of efficient sensor

calibration systems" existed. EX1005, ¶¶[0001], [0003]; EX1003, ¶55. Choe calibrates a camera for lane detection by receiving still or video images over a network; for video, the "images" are frames, as shown in Figures 6A–6B. EX1005, ¶¶[0018], [0022], FIGs. 6A–6B; ; EX1003, ¶56.



EX1005, FIG. 6A.



EX1005, FIG. 6B.

After receiving video frames, Choe determines a horizon line "based on the camera's hardware settings" or via "predictive models," EX1005, ¶¶[0015], [0030], and identifies horizon and lane lines through a perception process. *Id*., ¶¶[0016], [0056], [0059], FIG. 7; EX1003, ¶57. Choe then superimposes the horizon and lane lines onto the video frames, rendering images with the overlays. *Id*., ¶¶[0016], [0050]; EX1003, ¶58.



EX1006, FIG. 6A (excerpt, annotated); EX1005, FIG. 6A.

Choe overlays horizon/lane lines and displays them in a "visualizer" UI, EX1005, ¶[0014; the user aligns the overlays via keyboard/joystick/voice, e.g., arrow keys to adjust line position. EX1005, ¶¶[0014]–[0015], [0055]–[0057]; EX1003, ¶59. Choe then recalculates camera parameters, e.g., pitch from the difference between the initial and updated horizon (and similarly yaw/roll), and transmits the updated parameters back to the vehicle, EX1005, ¶¶[0014]–[0016], [0057]; EX1003, ¶¶60-61.

B. Westmacot

Westmacot addresses generating training data for a machine-learning road detector by automating/semi-automating road-image annotation, replacing labor-intensive manual labeling. EX1008, pp. 1-3; EX1003, ¶71. Westmacot annotates road structure (road surface, lanes) from frames captured as a vehicle drives,

EX1008, pp. 1, 22, detects a horizon and lane lines, and computes camera "reference parameters" (e.g., height, orientation), *id*., pp. 24, 44, 3233; EX1003, ¶¶72-80. A "human fixer" then refines results, adjusts/adds lines and edits parameters (e.g., height/crop), via the annotation UI shown in Figures 18–24. EX1008, pp. 19, 34; EX1003, ¶¶72-80. Examples of the user interface are shown below.



Fig. 18

EX1008, FIG. 18.

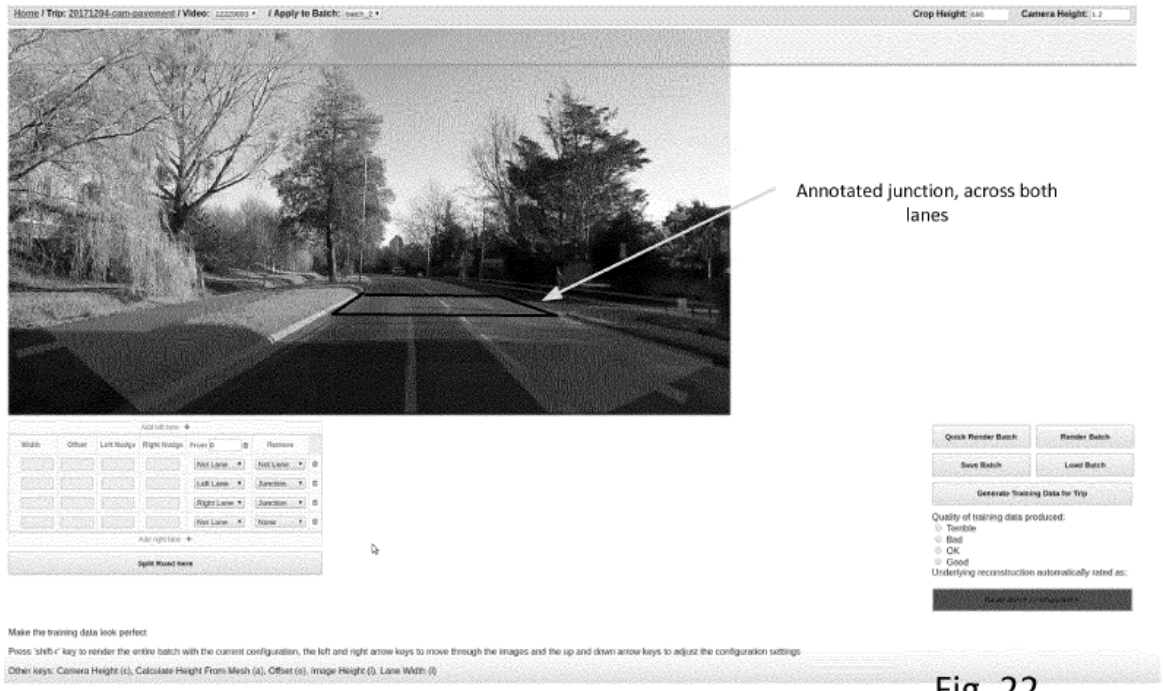


Fig. 22

EX1008, FIG. 22.

Westmacot determines camera parameters directly from captured images. EX1008, pp. 3, 6-7. Westmacot explains that "one or more parameters of the image capture device … may be derived from (only) the image data captured by the image capture device." *Id*., p. 3. For example, the expected road structure may be determined from a road normal vector, defined as "the surface normal of the road relative to the camera." *Id*., pp. 5, 25, which can be estimated from the image as:

$$\mathbf{n} = \frac{1}{\sum_{i=2}^{N-2} \|\mathbf{n}_i\|} \sum_{i=2}^{N-2} \mathbf{n}_i$$

$$\mathbf{n}_i = \mathbf{R}_i^{-1}(\mathbf{m}_{i-1,i} \otimes \mathbf{m}_{i,i+1}),$$

Id., pp. 5, 24, 26-27, 39-40, FIGs. 8-9, 16, 27 (additional details on the road plane normal).

Westmacot's Figure 26, shown below, illustrates the camera height (h) and road normal (n) in relation to the determined lane lines.

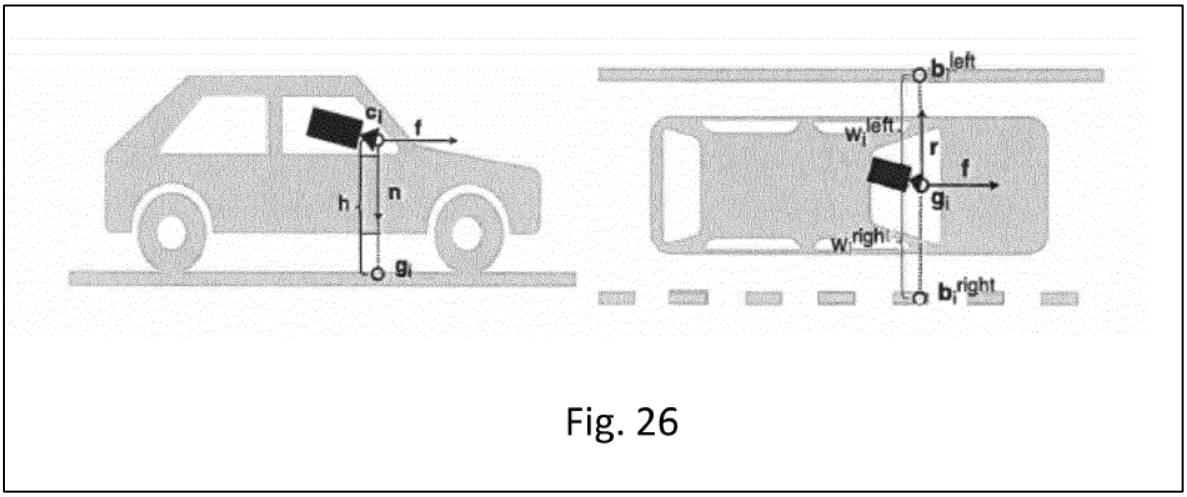


Fig. 26

EX1008, FIG. 26.

C. Davies

Davies is directed to "adjusting an image, e.g., an image horizon, for a vehicle mounted camera." EX1007, ¶[0002]; EX1003, ¶62. It describes a system that adjusts images "in response to at least one vehicle mounted sensor," EX1007,

¶[0006], with adjustments performed automatically using "telemetry of a vehicle from a plurality of sensors," *id*., ¶[0007]; EX1003, ¶¶63-67.

Davies further discloses that these adjustments may be performed onboard or remotely over a network. EX1007, ¶¶[0042]–[0044], [0059]; EX1003, ¶68. As shown in Figure 8, the vehicle's server 101 communicates with a third-party server 103 via network 106, where the server 103 may "include an application server providing applications and/or computer executable code implementing any of the interfaces/methodologies disclosed."EX1007, ¶[0064], FIG. 8; EX1003, ¶¶69-70.

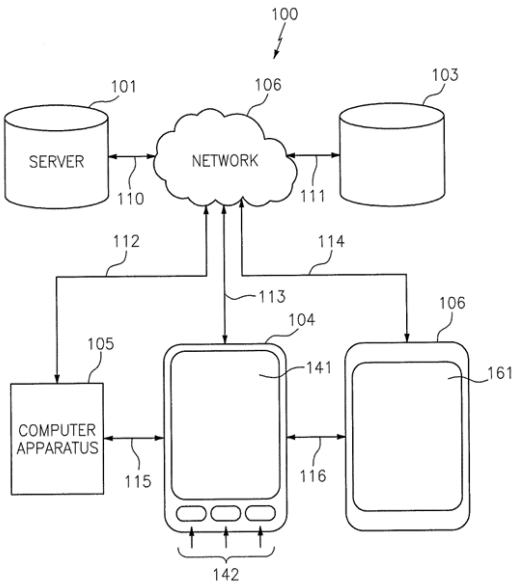


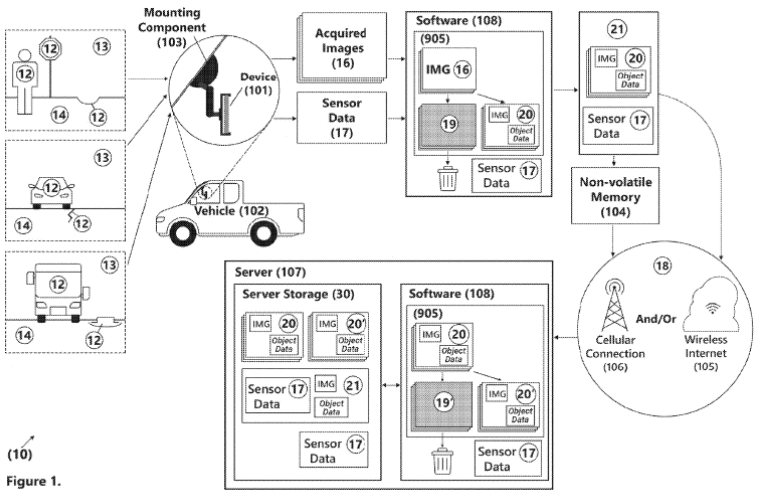
FIG. 8

EX1007, FIG. 8.

D. Tal

Tal discloses systems for image processing that "captur[e] digital images containing objects of interest … [and] transmit[] object data … over a communications network to a server … located remotely from the system." EX1009, ¶[0012]; EX1003, ¶¶81, 83. Using neural networks, Tal identifies "landmarks" such as "lane markings 12 and edge 12 on a road." *Id.*, ¶[0116].

As shown in Figure 1, digital images may be captured by "one or more imaging devices 101 … preferably with network communication capabilities to a communications network 18." EX1009, ¶[0027], FIG. 1; EX1003, ¶82. The images are transmitted to a remote server for further processing, and Tal explains it is "advantageous … to send object data 21 by the device 101 to the server 107 … in order to take advantage of the image 16 processing and acquisition efficiencies." EX1009, ¶[0027]; EX1003, ¶84-87.



EX1009, FIG. 1.

E. Kuehnle

Kuehnle discloses a system "for online calibration of a video system … in connection with an image-based road characterization … for detecting roadway scenes in vehicles." EX1010, 1; EX1003, ¶88. To determine camera orientation, Kuehnle detects vanishing points using a "long term average vanishing point" computed through time-filtering methods. EX1010, 3; EX1003, ¶¶88-89. These time-filtering methods consider both vehicle speed and elapsed time. EX1010, 11–13; EX1003, ¶89. Kuehnle specifies that extrapolation of vanishing points "require[s] that the vehicle is moving with at least a certain speed, so that low speed maneuvering … is not taking place." EX1010, 11–12; EX1003, ¶89. The system "calculates a long term average vanishing point location … from a sequence of images," from which "the static yaw and pitch angle of the camera is deduced." EX1010, 3, 13; EX1003, ¶90.

V. GROUND 1: CHOE RENDERS OBVIOUS CLAIMS 1-3, 5-10, 12-16 AND 18-20.

Choe teaches the full '276 calibration workflow: capturing roadway images; detecting horizon/lane lines; superimposing and user-updating those lines; computing pitch/yaw/roll from the updated lines; and using a client–server architecture that performs calibration off-board and returns results to the vehicle, thus rendering independent claims 1, 8, and 15 (and dependents) obvious. See

EX1005, ¶¶[0014]–[0018], [0025], [0031], [0044]–[0053], [0059]–[0060], FIGs. 1, 6A–6B, 7; EX1003, ¶¶100, 102.

To the extent Patent Owner disputes explicit transmission of the overlaid image and receipt of the updated line, a POSA would implement the disclosed user-adjustment at Choe's server: Choe already provides the line-editing interface and the ADV-network, so sending the overlaid image to the server for review and returning the new line (second position) is a routine within-reference allocation of disclosed functions with predictable results. *See* EX1005, ¶¶[0016]–[0018], [0025], [0031], [0044]–[0053]; EX1003, ¶101.

A. Independent Claim 1

1. [1Pre] "A method comprising"

Choe discloses this element, describing "processes or *methods* … [that] may be performed by processing logic that comprises hardware (e.g., circuitry, dedicated logic, etc.), software (e.g., embodied on a non-transitory computer readable medium), or a combination of both." EX1005, ¶[0064]: EX1003, ¶¶104-105.

2. [1A] "receiving an image of a roadway recorded by a camera device installed within a vehicle"

Choe discloses this element. EX1003, ¶¶106-111. Choe's ADV includes cameras 211 "to capture images of the environment surrounding the autonomous vehicle," EX1005, ¶¶[0020]–[0022], FIG. 2, "in response to a first image captured

by a camera of an ADV, a horizon line is determined … [and] one or more lane lines are determined" via perception, EX1005, ¶[0015], and sensor data includes "an image … processed by perception module 302 … determining a horizon line representing a vanishing point of a road," EX1005, ¶¶[0045], [0048]. Claim 1 likewise recites determining a horizon line from "a first image … [2D] view from the viewpoint of the ADV." Accordingly, Choe teaches "receiving an image of a roadway recorded by a camera device installed within a vehicle." EX1003, ¶¶106-111.

3. [1B] "detecting a horizon line in the image"

Choe discloses this element in two ways. EX1003, ¶¶112–122.

First, Choe discloses detecting a horizon line from an image captured by the ADV's camera: "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." EX1005, ¶¶[0015], [0048]. Perception module 302 performs "image processing … determining a horizon line representing a vanishing point of a road." *Id*., ¶[0045]; EX1003, ¶114. Figure 6A illustrates horizon line 610 "determined and placed on a location within image 601 based on the hardware setting of the targeted camera." EX1005, ¶[0055].



EX1006, FIG. 6A (excerpt, annotated).

Process 700 similarly confirms that "processing logic determines a horizon line … based on a pitch angle of a camera that captured a first image." EX1005, ¶[0059]; EX1003, ¶116.

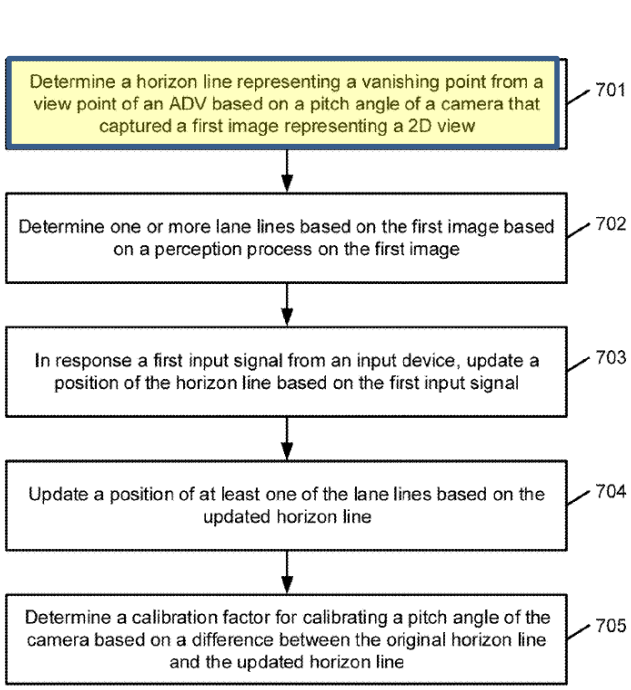


FIG. 7

EX1006, FIG. 7.

Second, Choe discloses detecting a horizon line using predictive models. Machine learning engine 122 generates predictive models 124, including calibration algorithms for use by ADVs. EX1005, ¶[0030], FIG. 1. A POSA would have understood these models predict horizon and lane lines, which may then be manually adjusted ("[t]he input signal may represent an incremental adjustment … of the horizon line"). *Id.*, ¶[0015]; EX1003, ¶¶118-119. After prediction and adjustment, the horizon line is re-identified and used by other modules: "sensor data is … processed by perception module 302 … determining a horizon line … [and] the result … can be utilized by other modules such as prediction module

303." EX1005, ¶[0045]; EX1003, ¶¶120-121.

Thus, Choe discloses "determining" a horizon line in an image captured by an in-vehicle camera, as recited in claim 1.

4. [1C] "overlaying a line on the image to generate an overlaid image"

In either form of horizon line detection described in Choe, direct detection by perception module 302 or predictive generation by models 124, Choe expressly discloses this element. EX1003, ¶¶123-126. Choe discloses superimposing the horizon line on the image (regardless of which form of detection), which a POSA would understand as "overlaying" the line. *Id.*, ¶124. For example, Choe explains that "the first image with the horizon line superimposed thereon is displayed … [and] updated … to give a visual feedback to a user." EX1005, ¶¶[0016], [0015], [0046], [0048], [0055]–[0059], FIGs. 6A–6B.

Choe's figures confirm this overlay. Figure 6A shows the horizon line superimposed on image 601 with "frame id: 1110," which Choe notes is "too low" and requires manual adjustment. EX1005, ¶[0057], FIG. 6A; EX1006, FIG. 6A. An image with the horizon line superimposed is, by definition, an "overlaid image." EX1003, ¶124.



EX1006, FIG. 6A (excerpt, annotated).

Accordingly, Choe discloses "overlaying a line on the image to generate an overlaid image" as recited in claim 1. EX1003, ¶126.

5. [1D] "transmitting the overlaid image to a computing device over a network"

Choe renders obvious this element. EX1003, ¶¶127–141. Choe expressly teaches a networked ADV-server architecture: the ADV "may be communicatively coupled to one or more servers 103–104 over a network 102," where "network 102 may be … [LAN, WAN such as the Internet, cellular, satellite], … wired or wireless," and "wireless communication system 112 … allow[s] communication … with servers 103–104 over network 102." EX1005 ¶¶[0018], [0025].

Choe also permits server-side calibration: "sensor calibration system 125 may be hosted by server 103 to calibrate sensors … offline based on the images and/or point clouds captured by the sensors," and "the calibration result can then

be uploaded onto the vehicle to be utilized online during the image processing." EX1005, ¶¶[0031], [0044]. At the same time, Choe teaches local/manual calibration in which "the first image with the horizon line superimposed thereon is displayed … [and] updated … to give a visual feedback to a user," with horizon and lane lines overlaid and adjusted. EX1005, ¶¶[0016], [0014]–[0017], [0049]–[0057], FIGs. 6A–6B. These disclosures confirm Choe's calibration imagery includes annotated overlays, not merely raw sensor data.

While Choe does not expressly state that these overlaid images are transmitted, a POSA would view that as a routine implementation in Choe's disclosed architecture. Choe already (i) generates annotated calibration views with horizon/lane lines "superimposed" and updated for user feedback, EX1005, ¶¶[0014]–[0017], [0049]–[0057], and (ii) performs calibration remotely: "sensor calibration system 125 may be hosted by server 103," with "the calibration result … uploaded onto the vehicle," EX1005, ¶¶[0031], [0044]), and over network 102 (LAN/WAN/Internet/cellular/satellite), EX1005, ¶[0018]. *See* EX1005, ¶[0025]. For remote calibration and user interaction to function as disclosed in Choe, the same processed, overlaid views must be sent so they can be rendered and adjusted at the remote terminal, merely allocating disclosed functionality across Choe's client–server design. *See* EX1005, ¶¶[0014]–[0017], [0018], [0025], [0031], [0044], [0049]–[0057]; EX1003 ¶¶128-132.

The motivations for transmission are straightforward. Sending the overlaid image ensures consistency by allowing remote personnel to see exactly what the vehicle displays, avoiding redundant detection and reducing server compute. Overlays already encode perception results and user adjustments, and annotated frames further support Choe's server-side training, drift monitoring, and audit functions. EX1005, ¶¶[0029]–[0031]; EX1003, ¶135. In human-in-the-loop scenarios, transmitting the same UI view used in-vehicle provides clear calibration context. EX1003, ¶138.

Implementation would require only routine design choices, either rasterizing overlays into a composite image or transmitting raw frames with overlay metadata, both obvious to try within Choe's disclosed networked framework and with a reasonable expectation of success. EX1005, ¶¶[0014]–[0017], [0049]–[0057], [0018], [0025]. Bandwidth concerns are immaterial, as overlays add negligible size and Choe already contemplates Internet, cellular, and offline uploads. EX1005, ¶¶[0018], [0031]. Nothing in Choe teaches away; transmitting annotated images is a predictable extension of its architecture. EX1003, ¶140.

6. [1E] receiving a modification of the line from the computing device, the modification comprising a new line at a second position;

Choe renders obvious this element. EX1003, ¶¶142-153.

Choe already teaches user-driven edits to overlaid roadway features: "in response to a second input signal … the position of at least one of the lane lines is modified … [and] a second calibration factor … is determined based on the modification," and "a user may modify a detected horizon line or lane line … [and] the calibration module may then compute a calibration parameter based on the modified line." EX1005, ¶¶[0017], [0052]–[0053]. These disclosures mean the system receives a modification yielding "a new line at a second position." At the same time, Choe also teaches a client–server framework: "Sensor calibration system 125 may be hosted by server 103 … based on the images … captured," with "the calibration result … uploaded onto the vehicle," and the ADV communicates via LAN/WAN/Internet/cellular/satellite. EX1005, ¶¶[0031], [0044], [0018], [0025]. Thus, Choe expressly contemplates the vehicle receiving updated calibration from a remote computing device (the server), i.e., receipt of the server-side modification to the line/line-based parameters.

In view of these teachings, a POSA would implement the same line-adjustment at the server: transmit the overlaid frame to server 103, have the server receive the user's edit, i.e., a new line at a second position, and return updated

calibration to the vehicle, because, in the Federal Circuit's words, selecting where to perform a disclosed function within a client–server system is a "simple design choice" among "two predictable choices." *Uber Techs., Inc. v. X One, Inc.*, 957 F.3d 1334, 1342–43 (Fed. Cir. 2020) ("a simple design choice," "two predictable choices," and "server-side plotting … was a design choice"); *see also CRFD Rsch., Inc. v. Matal*, 876 F.3d 1330, 1347 (Fed. Cir. 2017) (obvious where "there were only two predictable choices—perform the claimed method on the terminal or on the server—and both were within the technical grasp of the skilled artisan"); *ACCO Brands Corp. v. Fellowes, Inc*., 813 F.3d 1361, 1370 (Fed. Cir. 2016) ("The ordinary artisan would then be left with two design choices … [e]ach of these two design choices is an obvious combination of prior-art elements.").

Courts allow "common sense" to supply the rationale for straightforward implementation details. *See Perfect Web Techs., Inc. v. InfoUSA, Inc.*, 587 F.3d 1324, 1329–30 (Fed. Cir. 2009) ("Common sense has long been recognized to inform the analysis of obviousness if explained with sufficient reasoning."); *Wyers v. Master Lock Co*., 616 F.3d 1231, 1240–41, 1245–46 (Fed. Cir. 2010) (obviousness may rest on "logic, judgment, and common sense," and some implementations are "simply a matter of common sense"). Placing Choe's disclosed line-editing on the server is a predictable, routine variant that improves remote verification (letting reviewers "see what the vehicle sees") and reduces in-

vehicle load, without changing Choe's operation. EX1005, ¶¶[0015]–[0017], [0049]–[0053], [0018], [0025], [0031], [0044]; EX1003, ¶¶145-152.

7. [1F] computing a camera parameter based on the new line

Choe discloses this element. EX1003, ¶¶154–162. This step does not limit where computation occurs. Whether performed locally or remotely, Choe discloses the element. *Id*. Choe teaches that "in response to a first image … a horizon line is determined … [and] in response to a first input signal … the position of the horizon line is updated … [and] a first calibration factor … is determined … based on a difference between the initial horizon line and the updated horizon line." EX1005, ¶¶[0015], [0049]. Likewise, "in response to a second input signal … the position of at least one of the lane lines is modified … [and] a second calibration factor … for … yaw … [and] in response to a third signal … a third calibration factor … for … roll." EX1005, ¶¶[0017], [0052]–[0053].

Choe's process flow confirms: "At block 705, processing logic determines a calibration factor … based on the difference between the initial horizon line and the updated horizon line." EX1005, ¶[0059], FIG. 7.

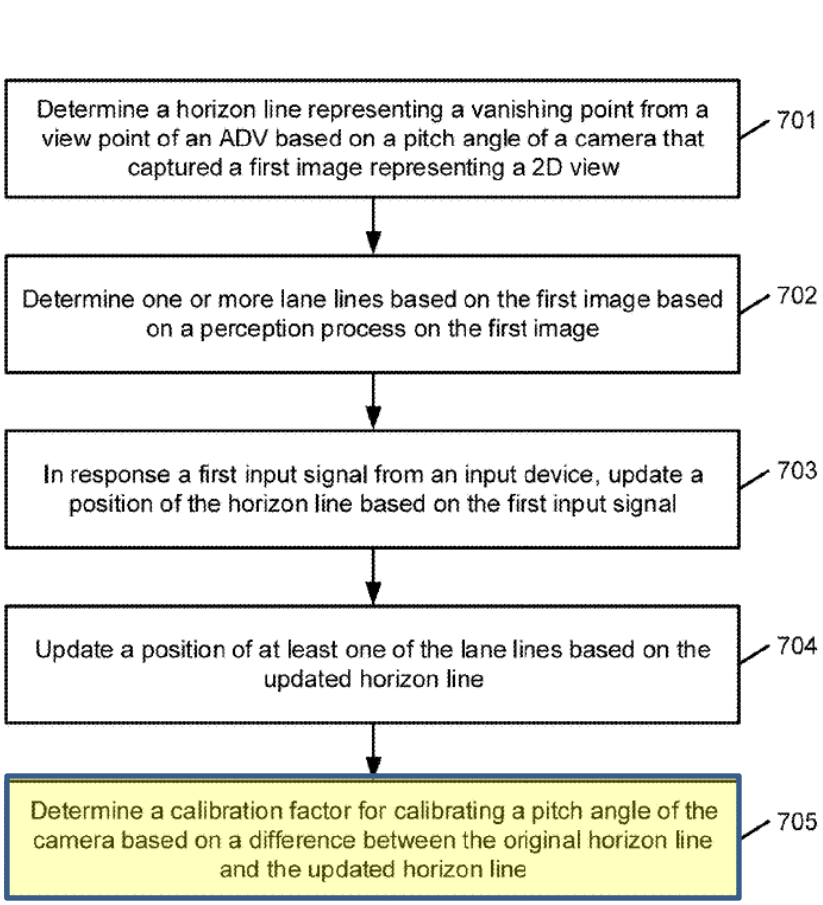


FIG. 7

EX1005, FIG. 7 (annotated).

A POSA would recognize these "calibration factors" as camera parameters, described as correction values used to calibrate the pitch, yaw, and roll of the camera hardware. EX1005, ¶¶[0015], [0017], [0049]–[0053], [0059]. Choe further explains that these values may be computed locally by calibration system 308 or remotely by server-based calibration system 125. *Id.*; EX1003, ¶¶159-161.

In short, Choe explicitly discloses that each time a line is updated to a new position, calibration factors (*i.e.*, camera parameters) for pitch, yaw, and roll are

computed from that line. EX1003, ¶162. Thus, Choe discloses element 1[F].

8. [1G] transmitting data representing the camera parameter to the camera device.

Choe discloses this element. EX1003, ¶¶163-169. Regardless of whether the claim is satisfied by (i) in-vehicle delivery (e.g., an in-vehicle calibration module sending updated pitch/yaw/roll over the vehicle bus) or (ii) off-board computation with parameters returned and forwarded to the camera, Choe teaches both.

Choe confirms that calibration parameters are not merely generated but fed back. Choe describes feedback of computed calibration values to adjust the sensor: "the calibration factors can be utilized by software applications such as a perception module to adjust the image processing parameters such as pitch, yaw, and roll angles to compensate the hardware during the image processing." EX1005, ¶[0017]. It further states that calibration factor calculator 406 "feeds [calibration factors] back to perception module 302 or other software modules such that the software modules can perform proper adjustment … to compensate the hardware settings of the sensors." EX1005, ¶[0047]. Choe also teaches remote/offline calibration with return of results: "sensor calibration system 125 may be hosted by server 103 … offline based on the images … captured by the sensors," and "the calibration result can then be uploaded onto the vehicle to be utilized online." EX1005, ¶¶[0031], [0044].

Accordingly, Choe discloses not only generating calibration parameters from updated line positions but also transmitting those parameters, whether locally via modules 406/308 or remotely from server 103, back to the camera device to update its operation. EX1003, ¶¶166-168. Thus, Choe discloses element 1[G].

B. Claim 2: "the camera parameter comprises one of camera height, viewing angle, and road plane normal."

Choe renders obvious camera parameters corresponding to camera height, viewing angle, and road-plane normal. EX1003 ¶¶170–181. After overlaying horizon and lane lines, Choe teaches that "in response to a first input signal (e.g., an up or down arrow key) … a position of the horizon line is updated … [and] a first calibration factor … for calibrating a pitch angle of the camera [is] determined." EX1005 ¶¶[0015], [0049]. Similarly, second and third inputs yield calibration factors for yaw and roll. EX1005 ¶¶[0017], [0052]–[0053].

Mapping is straightforward: the claimed "viewing angle" corresponds to Choe's pitch from horizon displacement; the "road-plane normal" follows from the combined pitch, yaw, and roll orientation; and although Choe does not state "camera height," a POSA would recognize that pitch calibration depends on camera mounting height, making it obvious in the same geometry. EX1003 ¶174.

Accordingly, Choe's computed pitch, yaw, and roll are interchangeable formulations of the claimed camera parameters, satisfying claim 2. EX1005 ¶¶[0015], [0017], [0049]–[0053]; EX1003 ¶¶180–181.

C. Claim 3: "receiving the image of the roadway comprises receiving an image of the roadway while the vehicle is moving."

Choe discloses this element. EX1003, ¶¶182-186. Choe expressly states: "The above processes can be utilized online while the vehicle is driving and real-time images are captured to calibrate the sensors in real-time," EX1005, ¶[0017]. Choe further shows horizon/lane updates and corresponding pitch/yaw/roll computation during operation, EX1005, ¶¶[0052]–[0054], [0059] (Block 705), and confirms that "while autonomous vehicle 101 is moving along the route," the system obtains real-time environment data from sensors, EX1005, ¶[0028]. Accordingly, Choe teaches receiving roadway images and calibrating while the vehicle is moving, satisfying claim 3. EX1005, ¶[0028]; EX1003, ¶186. Accordingly, Choe discloses claim 3.

D. Claim 5: "transmitting the overlaid image to the computing device over the network comprises transmitting the overlaid image to a web-based application."

Choe renders obvious claim 5. EX1003, ¶¶187-196. Choe teaches a networked architecture where autonomous vehicles communicate with remote servers over networks such as the Internet, cellular, or satellite: "Network configuration 100 includes autonomous vehicle 101 … coupled to one or more servers 103–104 over network 102," which "may be … the Internet, a cellular network, [or] a satellite network." EX1005 ¶[0018]. "Wireless communication system 112 … can wirelessly communicate with … servers 103–104 over network

102." EX1005 ¶[0025]. Calibration may occur remotely, with results uploaded back to vehicles: "Sensor calibration system 125 may be hosted by server 103 … [and] the calibration parameters … uploaded to the corresponding vehicles." EX1005 ¶¶[0031], [0044].

A POSA would have understood these disclosures render it obvious to transmit overlaid images to a web-based application. EX1003, ¶¶191-195. Choe already teaches generating overlaid images locally, EX1005, ¶¶[0015]–[0017], [0055]–[0057], transmitting sensor imagery and calibration data over the Internet to servers, EX1005, ¶¶[0018], [0025], [0031], [0044], and server-hosted calibration functionality, EX1005, ¶¶[0031], [0044]. A "web-based application" is simply a conventional way to implement server-hosted processing and visualization over the Internet. EX1003, ¶191. Using a web-based interface would have been a predictable choice among finite design options (e.g., desktop client, native application, or web application), yielding no unexpected result.

The motivation would have been clear: a web-based application enables calibration personnel or servers to view and verify annotated images (horizon and lane line overlays) without requiring specialized local software, supporting Choe's goals of server-hosted calibration, fleet-wide analytics, and human-in-the-loop verification. EX1005, ¶¶[0029]–[0031], [0044].

In *KSR*, the Supreme Court held that when there are "a finite number of identified, predictable solutions," pursuing a known option is likely the product of ordinary skill and common sense, not innovation. *KSR*, 550 U.S. at 421. Selecting a web-based interface instead of a desktop or native client to implement Choe's server-hosted calibration is precisely such a predictable design choice, no different than the application of "common sense" found obvious in *Perfect Web Techs.*, 587 F.3d at 1329–30; *see also Wyers*, 616 F.3d at 1239 (obviousness may rest on combining familiar elements "according to known methods" to yield predictable results); *Leapfrog Enters., Inc. v. Fisher-Price, Inc.*, 485 F.3d 1157, 1161 (Fed. Cir. 2007) (updating known systems with conventional technology is obvious).

Accordingly, transmitting the overlaid image to a web-based application is an obvious implementation of Choe's disclosed Internet-based calibration system. EX1003, ¶196.

E. Claim 6: "receiving the modification of the line from the computing device comprises receiving an identification of a new horizon line by a user of the computing device."

Choe renders obvious this element. EX1003, ¶¶197-204. Choe expressly teaches user-driven overlay adjustment—"in response to a second input … the position of at least one of the lane lines is modified … [and] a second calibration factor … is determined based on the modification." EX1005 ¶[0017]; see also ¶¶[0015], [0049]–[0053]. A POSA would view this as the user identifying a new

horizon-line position adopted for pitch/yaw/roll calibration. EX1003 ¶198.

Choe also discloses a client-server calibration setup where captured images are processed remotely and results uploaded back to the vehicle. EX1005 ¶¶[0018], [0025], [0031], [0044]. A POSA would naturally allocate the user-adjustment step to the server so the overlaid image is transmitted for review and the updated line returned—merely distributing known functions within Choe's architecture with predictable results. EX1003 ¶¶199–200.

Under *KSR*, this is straightforward: Choe teaches user-guided line setting with recalibration from the updated line, EX1005 ¶¶[0015], [0017], [0049]–[0053], [0055]–[0057], and networked ADV-server calibration. EX1005, ¶¶[0018], [0025], [0031], [0044]; EX1003, ¶¶210-211. Implementing the user-identification step at the disclosed server (sending the overlaid image for remote review and returning the modified line) is a permissible within-reference combination: "[c]ombining two embodiments disclosed adjacent to each other in a prior art patent does not require a leap of inventiveness." *Boston Scientific Scimed, Inc. v. Cordis Corp.*, 554 F.3d 982, 991 (Fed. Cir. 2009). More broadly, the Federal Circuit has confirmed that an obviousness determination may rest on a single prior-art reference when that reference discloses all elements; in that posture, the Board "was not required to make any finding regarding a motivation to combine." *Realtime Data, LLC v. Iancu*, No. 18-1154, slip op. at 8–9 (Fed. Cir. Jan. 10, 2019). Likewise, placing

Choe's already-disclosed edit operation on its already-disclosed server is a "simple design choice" between "two predictable choices" (server-side vs. device-side) that yields predictable results under *KSR. Uber Techs., Inc. v. X One*, Inc., 957 F.3d 1334, 1342–47 (Fed. Cir. 2020). Accordingly, Choe renders obvious claim 6. EX1003, ¶¶201-204.

F. Claim 7: "transmitting the data representing the camera parameter to the camera device further comprises recomputing the camera parameter based on the modification of the line."

Choe discloses this element. EX1003, ¶¶205-212. Choe explicitly discloses that each time a line is modified, the calibration system recomputes the corresponding camera parameter. EX1003, ¶206. For example: "in response to a first input signal … a position of the horizon line is updated … [and] a first calibration factor … is determined for calibrating a pitch angle of the camera based on a difference between the initial horizon line and the updated horizon line." EX1005, ¶¶[0015], [0049]. Likewise: "in response to a second input signal … the position of at least one of the lane lines is modified … [and] a second calibration factor for calibrating a yaw angle … [and] in response to a third signal … a third calibration factor for calibrating a roll angle …" EX1005, ¶¶[0017], [0052]–[0053].

Choe's Figure 7 confirms this iterative process: "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" EX1005, ¶[0059]. Each adjustment to the horizon or lane line triggers recomputation of the calibration factor. EX1003, ¶210.

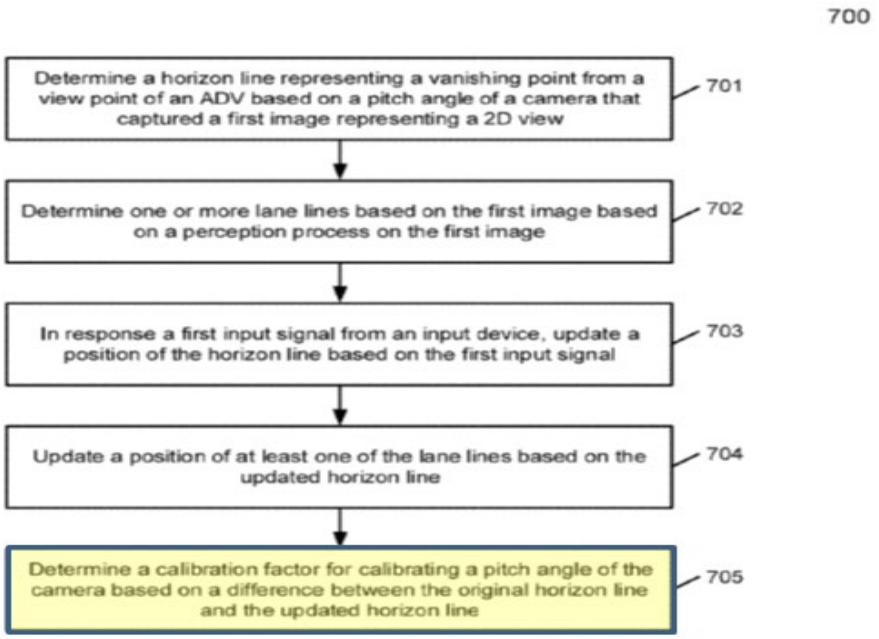


FIG. 7

EX1005, FIG. 7 (annotated).

Choe further explains that once these recalculated parameters are determined, they are transmitted back to the system to adjust the camera or applied by perception software to compensate the hardware: "The calibration factors can be utilized by software applications such as a perception module to adjust the image processing parameters such as pitch, yaw, and roll angles to compensate the hardware during the image processing". EX1005, ¶¶[0017], [0047], [0031], [0044]

(upload to vehicle); EX1003, ¶210. Accordingly, Choe discloses claim 7. EX1003, ¶212.

G. Claims 8-10, 12-16 18-20 are obvious over Choe.

Claims 8-10, 12-16 18-20 are likewise rendered obvious by Choe. EX1003, ¶¶213-219. These claims merely restate the substantive limitations of claims 1-3 and 6-7 in alternative statutory forms (computer-readable medium and device). Claim 8 recites the method of claim 1 in CRM form; claims 9-10 and 13-14 correspond to dependent claims 2-3 and 6-7; claim 15 mirrors claims 1 and 8 in device form; and claims 16 and 19-20 track the features of claims 3 and 13-14.

Choe discloses implementing the calibration functionality in both software and hardware, including perception module 302, calibration module 308, processors, and storage. EX1005, ¶¶[0032]-[0034], [0044], [0047]. Because claims 8-10, 12-16 18-20 add no substantive limitations beyond those already rendered obvious in claims 1-3 and 5-7, they too are rendered obvious by Choe. EX1003, ¶¶213-219.

VI. GROUND 2: CHOE IN VIEW OF DAVIES RENDERS OBVIOUS CLAIMS 1-3, 5-10, 12-16 AND 18-20.

Choe in view of Davies renders obvious claims 1-3, 5-10, 12-16 and 18-20.

A. A POSA would have been motivated to combine Choe with Davies.

1. Choe

As explained in Ground 1, Choe discloses a complete calibration workflow in which the ADV captures roadway images, determines and overlays horizon/lane lines, allows user adjustments to those lines, and recomputes pitch/yaw/roll from the updated geometry. EX1005, ¶¶[0015], [0017], [0049]–[0053], [0059]; EX1003, ¶220. It also teaches a client-server architecture where images are sent to a server-hosted calibration system and results are uploaded back to the vehicle. EX1005, ¶¶[0018], [0025], [0031], [0044]; EX1003, ¶220. A POSA would recognize that, for such remote calibration and user interaction to function, the same processed, annotated views used locally must be transmitted over the network so they can be rendered at the remote terminal. EX1003, ¶220. Choe suggests this arrangement but does not explicitly spell out the transport mechanics. *See* EX1005, ¶¶[0014]–[0017], [0031], [0044], [0049]–[0057]; EX1003, ¶¶220-221. Looking for routine implementation details, a POSA would consult complementary references, and Davies provides exactly that by expressly teaching wireless, bidirectional transmission of processed/annotated imagery and calibration data between a vehicle device and a remote server, with remote modifications returned to update

local parameters. EX1007, ¶¶[0059], [0063]–[0069], claim 24; EX1003, ¶¶222-223.

2. Davies

Davies squarely teaches remote transmission, operator adjustment, and return of processed outputs. It describes capturing 4K images onsite and "transport[ing] the captured images to an operations base ('OB') … e.g., a production truck away from the field," EX1007, ¶[0042]; at the OB, an operator uses a GUI "to navigate the full raster 4K image and maneuver the selective … extraction window," EX1007, ¶[0044]. The "output … is provided to a router … [and] can be taken live … or … ingested at a server … for later playout." EX1007, ¶[0045]. Davies further teaches distributing processing to address bandwidth: "some or all of image adjustment may be performed on the vehicle … particularly useful … in wireless transmission applications where a reduced data package can take advantage of bandwidth limitations," EX1007, ¶[0059], and confirms that an "adjusted image is transmitted via wireless protocol to an external computing device." EX1007, claim 24.

Taken together, these disclosures show a bidirectional, networked workflow in which processed/annotated imagery is sent to a remote computing device for human interaction and the adjusted output is returned for further processing, which

are implementation details that aligns directly with Choe's remote calibration framework.[1]

3. *KSR*

A POSA would have known how and why to combine Choe with Davies and would have had a reasonable expectation of success because the references disclose complementary elements that operate exactly as KSR predicts.

Choe discloses the full calibration loop: capture roadway images, overlay horizon/lane lines, accept user adjustments, and recompute pitch/yaw/roll and a

---

[1] To the extent Patent Owner contends that Davies's remote OB workflow and its optional on-vehicle processing are distinct embodiments, combining embodiments within a single reference is permissible where a POSA would do so. The Federal Circuit has held that "[c]ombining two embodiments disclosed adjacent to each other in a prior art patent does not require a leap of inventiveness." *Boston Scientific SciMed, Inc. v. Cordis Corp.*, 554 F.3d 982, 991 (Fed. Cir. 2009). A reference also "must be considered for everything that it teaches, not simply the described invention or a preferred embodiment." *In re Applied Materials, Inc.*, 692 F.3d 1289, 1298 (Fed. Cir. 2012) (citing *EWP Corp. v. Reliance Universal, Inc.*, 755 F.2d 898, 907 (Fed. Cir. 1985)).

client–server model where calibration "may be hosted by server 103" with results uploaded to the vehicle. EX1005, ¶¶[0015], [0017], [0049]–[0053], [0018], [0031], [0044]; EX1003, ¶¶226-228. Davies supplies the routine transport/UI layer: a server 101 with client devices 104-106 on network 106 that "may manipulate, share, transmit, and/or receive" system data; images are sent offsite to an operations base, adjusted via GUI 1114, and processed outputs routed back. EX1007, ¶¶[0066]-[0070], [0041]-[0045], [0047]; EX1003, ¶¶229-214. Davies also teaches bandwidth-aware split processing to "take advantage of bandwidth limitations." EX1007, ¶[0059]. A POSA would have found it obvious and understood *how* and *why* to combine Choe and Davies, transmitting and editing annotated calibration views off-vehicle represents a predictable combination of known elements according to their established functions with a reasonable expectation of success. EX1003, ¶234.

A POSA would have known *how* to combine these complementary teachings yields the claimed networked interaction in a predictable way: Choe's overlaid calibration image is transmitted to a remote device for user modification; the user identifies the corrected line at a second position; the updated geometry/parameters are returned and applied, which is nothing more than using known elements for their established functions. EX1005, ¶¶[0015], [0017], [0031], [0044], [0049]–[0053]; EX1007 ¶¶[0041]–[0047], [0066]–[0070]; EX1003, ¶¶233-234.

A POSA would have known *why* to combine Choe and Davies because the motivations are straightforward and well recognized: verification and consistency (send the same annotated view Choe already displays so remote reviewers "see what the vehicle sees," avoiding re-detection/divergence), human-in-the-loop oversight and fleet QA (server-hosted review across vehicles), computational offload (shift the adjustment/solve to the server Choe already contemplates), and bandwidth efficiency (Davies teaches reduced data packages to "take advantage of bandwidth limitations," implemented by transmitting lightweight overlay geometry/deltas and returning parameters). EX1005, ¶¶[0016], [0049]–[0057], [0018], [0031], [0044]; EX1007, ¶[0059]; EX1003, ¶¶235-245.

A POSA would have reasonably expected success in making this combination because both Choe and Davies address camera calibration in distributed environments, and each element functions as expected when combined. Nothing in either reference teaches away. Indeed, Choe expressly allows that "embodiments … may be modified without departing from the broader spirit and scope of the disclosure." EX1005, ¶[0066]; *see also DyStar Textilfarben GmbH v. C.H. Patrick Co*., 464 F.3d 1356, 1364 (Fed. Cir. 2006) (reference does not teach away absent criticism or discouragement); EX1003, ¶¶246-250.

B. Claim 1

Choe alone discloses or renders obvious every step of claim 1. Section V.A. To the extent Patent Owner contends that Choe alone does not render obvious elements 1[D] and 1[E], those features would have been obvious in view of Davies. Davies expressly describes wireless transmission of adjusted/processed images and image data to and from external computing devices, which further confirms the obviousness of performing elements 1[D]–1[G] over a network.. EX1007, ¶¶[0059], [0063]–[0069], claim 24; EX1003, ¶251.

1. 1[Pre]-1[C], 1[F], 1[G]

Choe discloses elements 1[Pre]-1[C], 1[F], 1[G]. Sections V.A.1-4, 7, 8. EX1003, ¶252. Choe discloses receiving roadway images from in-vehicle cameras, detecting and determining horizon/lane lines, and superimposing them on the image. EX1005, ¶¶[0014]–[0017], [0045]–[0046], [0049]–[0057], FIGs. 6A–6B. Choe also discloses computing computer calibration parameters from user-updated lines and transmitting data representing the camera parameter to the camera device. EX1005, ¶¶[0015], [0017], [0031], [0044], [0049]–[0053], [0059]–[0060].

2. 1[D]: "transmitting the overlaid image to a computing device over a network"

Choe in view of Davies renders obvious element 1[D]. EX1003, ¶253.

As discussed above, Choe renders obvious element 1[D]. But to the extent PO argues to the contrary, it certainly would have been obvious in view of Davies. EX1003, ¶255.

a. Choe

Choe describes a networked ADV coupled to servers over LAN/WAN/Internet/cellular links, EX1005, ¶¶[0018], [0025], server-hosted calibration "based on the images and/or point clouds captured by the sensors," *id*., ¶[0031], and uploading calibration results back to the vehicle, *id*., ¶[0044]. EX1003, ¶254. Choe also generates and updates overlaid images locally, e.g., horizon line 610 and lane lines 621–623, in real time. EX1005, ¶¶[0014]–[0017], [0049]–[0057], FIGs. 6A–6B; EX1003, ¶254.

b. Davies

Davies discloses a vision-processing/calibration system in which vehicle-captured images are sent to an external computing device for operator review and adjustment. For example, images are "transport[ed] … to an operations base ('OB') … e.g., a production truck away from the field," EX1007, ¶[0042], where a GUI lets an operator "navigate the full raster 4K image and maneuver the selective … extraction window," EX007, ¶[0044]. EX1003, ¶256. The "output … is

provided to a router … [and] can be taken live … or … ingested at a server … for later playout," EX1007, ¶[0045], allowing the adjusted imagery to be routed back into the system. EX1003, ¶256. Davies also teaches bandwidth-aware split processing: "some or all of image adjustment may be performed on the vehicle … particularly useful … in wireless transmission applications where a reduced data package can take advantage of bandwidth limitations." EX1007, ¶[0059]; *see also* EX1007, claim 24; EX1003, ¶256.

c. *KSR*

A POSA would have found it obvious to transmit overlaid (not just raw) images over Choe's existing network for several predictable, technical reasons: (i) **verification/consistency**: remote systems "see what the vehicle sees" without re-running overlay detection, reducing divergence; (ii) **computational efficiency**: overlays already encode local perception and user edits, so sending them avoids redundant processing on the server; and (iii) **fleet-level analytics and human-in-the-loop review**: central access to annotated frames enables troubleshooting, auditing, and standardization. EX1003, ¶¶257-262.

Choe supplies the overlay and networking framework, EX1005, ¶¶[0014]–[0017], [0049]–[0057], [0018], [0025], [0031], [0044]; Davies supplies the explicit bidirectional workflow, EX1007, ¶¶[0059], [0063]–[0069], claim 24. Implementing transmission of overlaid images would have been a predictable

implementation within the finite set of conventional options (e.g., rasterized composites or raw frames with overlay metadata), with a reasonable expectation of success. EX1003, ¶258.

Accordingly, Choe in view of Davies renders obvious element 1[D].

3. 1[E]: "receiving a modification of the line … comprising a new line at a second position"

Choe in view of Davies renders obvious element 1[E]. EX1003, ¶¶263-276. As discussed above, Choe alone renders obvious element 1[E]. But to the extent PO argues to the contrary, it certainly would have been obvious in view of Davies. EX1003, ¶266.

a. Choe

Choe teaches the substance of this step: after overlaying roadway features, the system accepts user input that "modif[ies]" the position of a line and then uses the modified line for calibration, e.g., "in response to a second input signal … the position of at least one of the lane lines is modified … [and] a second calibration factor … is determined based on the modification," and "a user may modify a detected horizon line or lane line … [and] the calibration module may then compute a calibration parameter based on the modified line." EX1005, ¶¶[0017], [0052]–[0053]; EX1003, ¶264. Choe also discloses a client-server calibration architecture in which "sensor calibration system 125 may be hosted by server 103 … based on the images … captured by the sensors," with "the calibration result …

uploaded onto the vehicle." EX1005, ¶¶[0031], [0044]; *see also* EX1005, ¶¶[0018], [0025]; EX1003, ¶265.

b. Davies

Davies provides the routine network and GUI framework for performing the same line-modification step remotely and returning results. Davies teaches sending captured images "to an operations base … away from the field," where an operator uses a GUI to "navigate the full raster 4K image" and the "output … [is] ingested at a server." EX1007 ¶¶[0042]–[0045]; EX1003 ¶267. Davies also discloses split on-vehicle/offsite processing for wireless transmission efficiency and a client-server network where users "manipulate, share, transmit, and/or receive" data. EX1007 ¶¶[0059], [0066]–[0070]; EX1003 ¶267.

c. *KSR*

Combining these complementary teachings in Choe and Davies yields the predictable result: the same user-driven line modification Choe performs locally is executed on the remote device (per Davies), and the calibration system receives that modified line, i.e., a "new line at a second position," over the disclosed network. This is a routine allocation of known functions across Choe's architecture, using Davies's conventional transport/GUI workflow, with obvious benefits (consistency, human-in-the-loop oversight, and bandwidth efficiency by sending lightweight line coordinates/overlays) and a reasonable expectation of

success. EX1005, ¶¶[0014]–[0017], [0031], [0044], [0049]–[0057]; EX1007, ¶¶[0042]–[0045], [0059], [0066]–[0070]; EX1003, ¶¶268-276.

Accordingly, Choe in view of Davies renders obvious element 1[E].

C. Claims 2-3 and 5-7

Claims 2-7 are likewise obvious over Choe in view of Davies for the same reasons set out under Ground 1, with Davies supplying any alleged gaps. Each dependent claim merely narrows the core overlay-modify-recompute workflow of claim 1 in ways already taught by Choe or that would have been obvious to a POSA once Davies's explicit two-way (vehicle ↔ external device) calibration framework is incorporated. EX1003, ¶277.

1. Claim 2 (camera height, viewing angle, road-plane normal).

Choe explicitly discloses claim 2. Section V.B; EX1003, ¶278.

Choe teaches computing pitch, yaw, and roll whenever a line is adjusted: pitch from horizon-line modifications and yaw/roll from lane-line changes. EX1005, ¶¶[0015], [0017], [0049], [0052]–[0053], [0059]. A POSA would have understood that these angles define the camera's viewing angle and, by straightforward application of projective geometry, the road-plane normal. EX1003, ¶278. Further, because Choe's pitch calculation is based on the horizon line relative to the roadway, a POSA would have recognized that this calculation directly corresponds to the camera's effective mounting height. EX1003, ¶278.

Claim 2 is therefore obvious. EX1003, ¶278.

2. Claim 3 (image received while the vehicle is moving).

Choe explicitly discloses claim 3. Section V.C; EX1003, ¶279.

Choe teaches online calibration "while the vehicle is driving and real-time images are captured," EX1005, ¶[0017], and continuous perception while the ADV "is moving along the route," *id*., ¶[0028]. Claim 3 is therefore obvious. EX1003, ¶279.

3. Claim 5 (transmit to a web-based application).

Claim 5 is obvious over Choe in view of Davies. EX1003, ¶¶280-285.

a. Choe

Choe does not expressly specify delivery via a web application, but it already contemplates the exact deployment model: an ADV communicates with remote servers "over a network 102," which "may be … the Internet," via a wireless communication system that talks to "servers 103–104 over network 102"; calibration "may be hosted by server 103 … based on the images … captured by the sensors," and "the calibration result can then be uploaded onto the vehicle." EX1005, ¶¶[0018], [0025], [0031], [0044]; EX1003, ¶280. Those teachings align with how a web-based application operates in practice: Internet-connected servers hosting an interface that renders the same overlaid calibration views Choe generates. *See id.*, ¶¶[0014]–[0017], [0049]–[0057]; EX1003, ¶280.

b. Davies

Davies reinforces and operationalizes this remote UI pattern: images captured on the vehicle are "transport[ed] … to an operations base," where an operator, via a GUI, "navigate[s] the full raster 4K image and maneuver[s] the selective … extraction window," with the processed output "provided to a router … [and] ingested at a server" for continued use; processing can be split between on-vehicle and remote processors "particularly useful … in wireless transmission applications," and modified data can be returned to the vehicle, EX1007 ¶¶[0042]–[0045], [0059], [0063]. That is functionally equivalent to hosting the calibration interface as a web-based application that receives and displays the overlaid image, enables adjustments, and returns updates. EX1003, ¶¶283-284.

c. *KSR*

A POSITA would have been motivated to implement the remote interface as a browser-based web application, one of a few routine alternatives (web client vs. native client), to integrate Choe's Internet-connected, server-hosted calibration framework with the remote GUI workflow taught by Davies. Doing so would simply substitute a known interface technology for the same client-server exchange, using Choe's existing network infrastructure to transmit calibration imagery and parameters. The combination would predictably allow remote operators to view and adjust overlaid calibration images through a standard

browser, with a reasonable expectation of success given that web visualization, HTTP streaming, and remote GUI editing were mature and commonly integrated technologies. *See* EX1005, ¶¶[0018], [0025], [0031], [0044]; EX1007, ¶¶[0042]–[0045], [0059], [0063]; EX1003, ¶284-285.

4. Claim 6 (receiving a user identification of a new horizon line).

Claim 6 is obvious over Choe in view of Davies. EX1003, ¶286.

a. Choe

Choe teaches user-driven horizon-line setting. Choe superimposes a horizon line on the captured frame and permits the user to adjust it via keys/joystick/voice, with calibration recomputed from the updated geometry: "in response to a first input signal (e.g., an up or down arrow key) … a position of the horizon line is updated … [and] a first calibration factor … is determined … based on a difference between the initial horizon line and the updated horizon line," and "a user may modify a detected horizon line or lane line by providing input through an input device … the calibration module may then compute a calibration parameter based on the modified line." EX1005, ¶¶[0015], [0052]–[0053]; *see also* EX1005, ¶¶[0017], [0049]; EX1003, ¶286. Functionally, the user's adjustment is an identification of the new horizon line position that the system receives and uses for calibration.

b. Davies

Davies supplies the networked, remote UI context in which the same user identification is performed off-vehicle and returned to the system. It discloses transporting captured images to an external operations base where an operator, via a GUI, "navigate[s] the full raster 4K image and maneuver[s] the selective … extraction window," with "processing … either offsite … or onsite," and the adjusted outputs routed back into the system; processing can be split between on-vehicle and remote processors "particularly useful … in wireless transmission applications." EX1007, ¶¶[0042]–[0045], [0059]; EX1003, ¶286.

c. KSR

A POSA would therefore implement Choe's line-adjustment step on the remote GUI so the external user identifies a new horizon line, transmits that identification over the network, and the vehicle/server uses it to recompute calibration, which is nothing more than applying known elements (Choe's user-set horizon line and Davies's remote, bidirectional GUI workflow) according to their established functions with predictable results. EX1005, ¶¶[0015], [0017], [0049], [0052]–[0053]; EX1007, ¶¶[0042]–[0045], [0059]; Section VI.A; EX1003, ¶286. Claim 6 is therefore obvious. EX1003, ¶286.

5. Claim 7 (recomputing the camera parameter based on the modification).

Claim 7 is obvious over Choe in view of Davies. EX1003, ¶287.

a. Choe

Choe already does exactly this: every line adjustment triggers recalculation of the relevant parameter (pitch/yaw/roll). EX1005, ¶¶[0015], [0017], [0049], [0052]–[0053], [0059]; EX1003, ¶287.

If Patent Owner argues that Choe does not describe performing this recomputation remotely, Davies fills this gap.

b. Davies

Davies reinforces and extends this by making clear that the adjustment/recalculation loop can run off-vehicle: captured images are transported to an offsite operations base where an operator manipulates them via a GUI, and "the output … is provided to a router … [and] can be taken live … or … ingested at a server … for later playout," with processing "either offsite … or onsite," and with processing split between on-vehicle and external processors "particularly useful … in wireless transmission applications." EX1007, ¶¶[0042]–[0045], [0059]; EX1003, ¶288.

c. *KSR*

A POSA would understand that Choe's same recalculation step (pitch/yaw/roll from the updated line) can be performed at the remote node and the

resulting adjusted data returned. Under KSR, combining Choe's explicit "modify line → recompute parameter" loop with Davies's routine offboard processing and return path is a predictable use of known elements. Whether the user moves the line locally or on a remote GUI, the system recomputes the same camera parameters from the new line position and applies them, which is an expected, straightforward implementation with a reasonable expectation of success. EX1005, ¶¶[0015], [0017], [0049], [0052]–[0053]; EX1007, ¶¶[0042]–[0045], [0059]; EX1003, ¶288.

A POSA would have been motivated to combine these complementary teachings. EX1003, ¶289. By the priority date, distributing computation between onboard and server-side processors was a standard design in autonomous driving systems to conserve vehicle resources and ensure consistency across fleets, as supported by contemporaneous art. Thus, applying Davies's routine remote-processing framework to Choe's recalculation workflow would have been the predictable use of known techniques, with a reasonable expectation of success. EX1003, ¶288. Claim 7 is therefore obvious. *Id.*

D. Claims 8-10, 12-16, and 18-20

Claims 8–10, 12–16, and 18–20 are likewise obvious over Choe in view of Davies. EX1003 ¶290. These claims merely restate the substantive limitations of claims 1–3 and 5–7 in computer-readable medium or device form.

Choe already discloses implementing the calibration functionality in software and hardware through its perception module 302, calibration module 308, processors, and storage, EX1005, ¶¶[0032]–[0034], [0044], [0047]. Davies reinforces that such calibration processes may be performed locally or remotely by an external computing device. EX1007, ¶¶[0059], [0063]–[0069]. Because claims 8-10, 12-16, and 18-20 add nothing of substance beyond the limitations already rendered obvious in claims 1-3 and 5-7, they too are unpatentable over Choe in view of Davies. EX1003, ¶290.

VII. GROUNDS 3A/3B: CHOE IN VIEW OF KUEHNLE (GROUND 3A) AND CHOE IN VIEW OF DAVIES AND FURTHER IN VIEW OF KUEHNLE (GROUND 3B) RENDER OBVIOUS CLAIMS 4, 11, AND 17.

Grounds 1 and 2 established the obviousness of the independent claims (Choe alone; Choe in view of Davies). Building on that foundation, Grounds 3A and 3B show dependent claims 4, 11, and 17 are likewise unpatentable: Choe in view of Kuehnle (Ground 3A) and Choe in view of Davies, further in view of Kuehnle (Ground 3B), render these claims obvious.

A. Claims 4, 11, and 17: "receiving the image of the roadway while the vehicle is moving comprises detecting that the vehicle is traveling above a pre-defined speed or for a pre-defined duration." (Grounds 3A/3B)

1. Choe (Ground 3A) and Choe in view of Davies (Ground 3B)

Choe in view of Kuehnle (Ground 3A) and Choe in view of Davies and further in view of Kuehnle (Ground 3B) render obvious claims 4, 11, and 17. EX1003, ¶¶291-301.

For both grounds, Choe already performs calibration "online while the vehicle is driving and real-time images are captured." EX1005 ¶[0017]. While driving, it collects "real-time … data detected … by sensor system 115" and logs "speeds, accelerations, [and] decelerations." EX1005 ¶¶[0028]–[0029]. Although Choe suggests use of speed and time telemetry, it does not explicitly link minimum speed or elapsed duration to triggering calibration updates. EX1003 ¶292.

Similarly, under Ground 3B, Davies ties calibration and image adjustment to vehicle dynamics, using telemetry such as "gyro data, vehicle angle, attitude, altitude, speed, acceleration, traction, [and] navigational data" to decide when updates occur. EX1007 ¶[0031]; EX1003 ¶293. Davies also bases calibration on multiple frames, describing "receiving image data … [and] data from … vehicle mounted sensor[s]" to refine the image horizon over time. EX1007 ¶[0057]; EX1003 ¶293. Yet Davies likewise does not explicitly condition calibration on speed or time thresholds.

A POSA considering Choe alone (Ground 3A) or Choe in view of Davies (Ground 3B) would have found it obvious to look to Kuehnle, which expressly requires calibration only when the vehicle moves "with at least a certain speed" and applies "time-filtering methods from a sequence of images." EX1010, 3, 11–13. These safeguards address the same issue recognized in Choe and Davies: ensuring calibration updates occur only under stable, representative driving conditions. EX1003 ¶¶294–301.

2. Kuehnle

Kuehnle discloses an online calibration method for vehicle-mounted video systems that determines camera pitch and yaw from vanishing points and requires calibration only when the vehicle travels above a set speed or for a defined duration. It states that calibration "require[s] that the vehicle is moving with at least a certain speed … so that low speed maneuvering … is not taking place," and that the system evaluates "the current speed data of the vehicle." EX1010, 11–12; EX1003 ¶294.

Kuehnle also applies a temporal filter, computing a "long term average vanishing point … from a sequence of images" and deducing yaw and pitch from that averaged data, with weighting that "decreases with time (as in a recursive averaging filter)." EX1010, 3, 13; EX1003 ¶295.

Accordingly, Kuehnle teaches both a predefined speed threshold and duration requirement, satisfying the claimed condition for calibration only when the vehicle is moving above a set speed or for a set time.

3. *KSR*

Under KSR (for both grounds), a POSA would have found it obvious to add Kuehnle's speed and duration conditions as a "combination of familiar elements according to their established functions," with only predictable options such as setting a minimum vehicle speed or elapsed duration. KSR, 550 U.S. at 421; EX1003 ¶¶297–298. Choe already provides the necessary telemetry, and Kuehnle shows how to use it—by conditioning calibration updates on predefined thresholds. Choosing a threshold (e.g., 5 km/h or 10 s) is a routine engineering step to avoid recalibration during idle or parking and ensure stable forward motion.

For Ground 3B, applying Kuehnle's time-filtering to Choe/Davies's multi-frame refinement is simply using known averaging to suppress transient noise. EX1003 ¶299.

Across both grounds, a POSA would expect success because each reference addresses the same problem (accurate in-motion camera calibration) and Kuehnle's thresholds mitigate the transient risks noted in Choe and Davies. The combination applies known elements in predictable ways, rendering claims 4, 11, and 17

obvious in view of Choe + Kuehnle or Choe + Davies + Kuehnle. EX1003 ¶¶300–301.

VIII. GROUND 4: WESTMACOT IN VIEW OF TAL RENDERS OBVIOUS CLAIMS 1-20.

A POSA would have found claims 1–20 obvious in view of Westmacot in combination with Tal. EX1003, ¶302.

A. A POSA would have been motivated to combine Westmacot and Tal.

1. Westmacot

Westmacot discloses a road-annotation and calibration system for autonomous vehicles. Vehicle-mounted cameras capture roadway images, from which geometric features like the horizon line are computed. An annotation interface overlays lane lines on the images and allows users to shift or add boundaries, with edits propagated through a 3D road model via model adaptation component 410. EX1008, 30–32, 34–35, FIGs. 3, 7, 15, 30–32; EX1003 ¶¶303–304. Parameter computation component 408 recalculates camera parameters—height (H), offset (S), forward point, horizon line, lane width (W), and center (C) and updates other modules accordingly. In sum, Westmacot teaches the same workflow as the challenged claims: image capture, overlay of horizon and lane lines, user modification, and recalculation of camera parameters.

2. Tal

Tal discloses a bidirectional, networked vision pipeline where both raw and processed imagery, including overlays, are exchanged between an in-vehicle device and a server. Selected image frames or unprocessed images may be transmitted in real time. EX1009 ¶[0027]; EX1003 ¶305. Connectivity includes Wi-Fi "to upload a large volume of data" and cellular links for continuous updates. EX1009 ¶¶[0064]–[0065]. The server "organizing, storing, processing and disseminating … object data 21" can store "image data 16, sensor data 17, [and] processed data 20," with processing split between client and server (Fig. 4). EX1009 ¶¶[0050]–[0051], [0066]. Overlays accompany the imagery: processed data "can be overlayed on the image … [and] flattened/merged," or "stored and sent separately." EX1009 ¶[0052]. Tal thus provides what Westmacot may lack: (i) network transmission of overlaid images, (ii) remote client-server processing and modification, and (iii) return transmission of processed results or parameters. EX1009 ¶¶[0027], [0050]–[0052], [0064]–[0066]; EX1003 ¶307.

3. *KSR*

Taken together, Westmacot and Tal teach the full scope of claims 1–20, and a POSA would have known *how* and *why* to combine with a reasonable expectation of success. EX1003 ¶308.

Westmacot captures roadway images, overlays horizon and lane lines, allows user modification, and computes camera parameters (e.g., height, offset, horizon line, lane width, center), but does not explicitly transmit these outputs beyond the local system. Given its goal of scalable training, remote access is implied. Tal fills any gap, teaching transmission of raw and processed images, including overlays and calibration data, between vehicles and servers for analysis and feedback. EX1009 ¶¶[0027], [0030], [0052], [0075], [0138]; EX1003 ¶308.

A POSA would combine these teachings to enable distributed processing, remote annotation, and fleet-wide calibration consistency. Tal's networked pipeline distributes Westmacot's annotated images and parameters for training and calibration, extends its UI for collaborative editing, and returns recalculated parameters (e.g., height, offset, horizon line) to the vehicle, closing the calibration loop. EX1008 FIGs. 6 (408), 25 (412); EX1009 ¶¶[0027], [0030], [0052], [0075], [0138]; EX1003 ¶309.

Westmacot's annotated outputs, confined to internal modules (e.g., 410, 416), are ideal for ML pipelines. EX1003, ¶311 Tal teaches transmitting these over wireless or cellular networks for centralized aggregation, distributed annotation, and calibration consistency—routine by the priority date. Integrating Tal's transmission framework would predictably enhance scalability and accuracy.

EX1009 ¶¶[0027], [0030], [0052], [0059], [0063], [0075], [0138]; EX1003 ¶¶310, 312.

Both references operate in the same field, autonomous-vehicle image processing and calibration, and complement each other: Westmacot generates the annotations; Tal enables their remote refinement. Their combination is the "predictable use of prior art elements according to their established functions." KSR, 550 U.S. at 417. Nothing teaches away, and a POSA would have every expectation of success. EX1003 ¶¶313-315.

B. Independent Claim 1

1. [1Pre] "A method comprising"

Westmacot discloses a method in autonomous vehicle contexts: vehicle-mounted cameras capture roadway frames; the system overlays horizon/lane lines, permits user adjustment, and updates camera/reference parameters based on those adjustments, mapping directly to the steps of claim 1. EX1003 ¶316; EX1008, Background; First/Second Aspects; FIGs. 3, 7, 15, 30–32.

2. [1A] "receiving an image of a roadway recorded by a camera device installed within a vehicle"

Westmacot discloses this element. EX1003, ¶317.

Westmacot explains that "training images are like the images that will be seen from cameras in the autonomous vehicle" and are annotated to mark the road surface or lane pixels. EX1008, 3. It further teaches "a method of annotating

frames of a time sequence of frames captured by at least one travelling vehicle" and "receiving a time sequence of images as captured by an image capture device of a travelling vehicle." EX1008, 1, 5, 9, 14, 68, 11; EX1003, ¶¶318-319. Westmacot also provides a concrete example: roadway videos "captured with a standard Nextbase 402G Professional dashcam … mounted on the inside of the car windscreen, roughly along the centre line of the vehicle." EX1008, 1, 5, 9, 14, 68, 11. Figure 3 shows "vehicle 300 … compris[ing] an image capture device 302" feeding images to processor 304 and memory 306. EX1008, FIGs. 3, 7, 15, 30–32; EX1003, ¶320.

Accordingly, Westmacot discloses "receiving an image of a roadway recorded by a camera device installed within a vehicle." EX1003, ¶321.

3. [1B] "detecting a horizon line in the image"

Westmacot discloses this element. EX1003, ¶¶322-330.

Westmacot explains that the road surface normal vector can be projected onto the image plane "to provide a line across the image that would match the horizon … (the 'horizon line' as that term is used herein)." EX1008, 44, FIG. 16; EX1003, ¶323. It further states that camera orientation is estimated from roadway geometry: "the vector difference xt+1 — xt-1 lies (approximately) parallel to the vehicle's longitudinal axis … to estimate the angular offset of the camera 302

relative to the vehicle 300," EX1008, 43, which is detection of the horizon line for pitch reference. EX1003, ¶324.

Westmacot also describes estimating "a 3D location of the camera 302 at the time that image was captured," EX1008, 37, which a POSA would understand requires horizon alignment to map 2D image coordinates to 3D world coordinates. EX1003, ¶¶325-329. Finally, Westmacot explains that horizon detection enables automatic lane annotation: "it is possible to fully automatically generate annotation of the lane … using only images captured from a low-cost and un-calibrated image capture device." ," EX1008, 48. Accordingly, Westmacot discloses "detecting a horizon line in the image." EX1003, ¶330.

4. [1C] "overlaying a line on the image to generate an overlaid image"

Westmacot discloses this element. EX1003, ¶¶331-337.

Westmacot teaches overlaying annotation lines on roadway images for user display. The rendering component "can render … captured images … overlaid with … annotation data A(n)," EX1008, 51, and Figure 6 shows lane boundaries R1 and R2 projected onto the image for user adjustment. The "annotation interface … provides a human annotator with the ability to view the frames … with the currently rendered lanes projected … [which] can be widened, narrowed and moved … Immediate feedback is provided via projection in the 2D camera view." EX1008, 35.

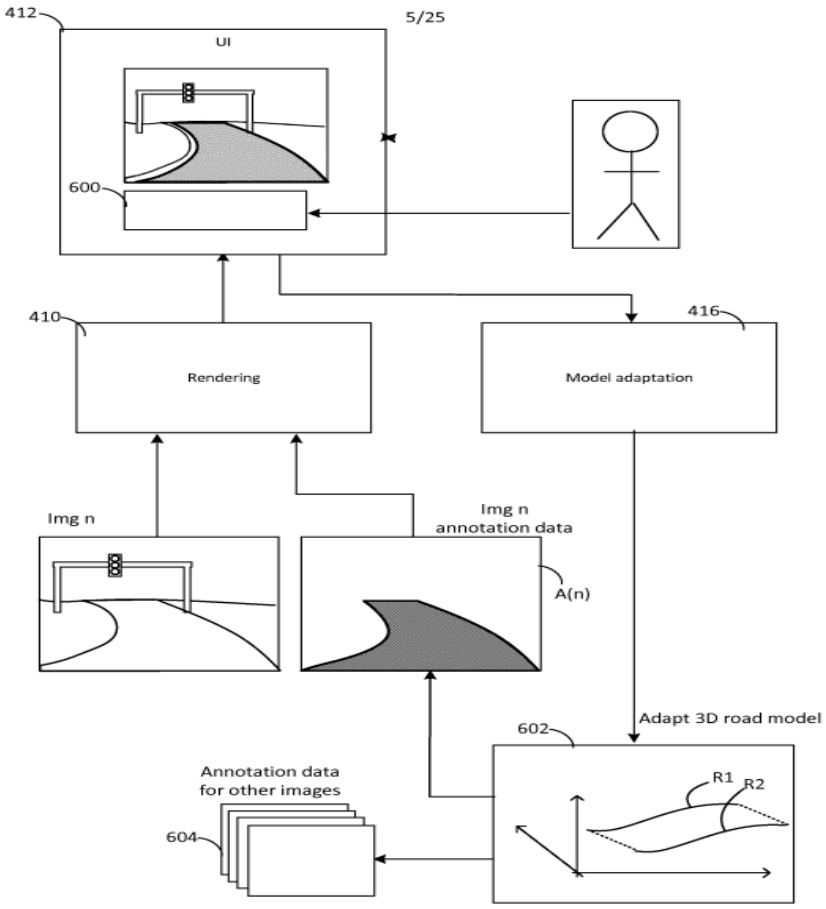


FIG. 6

EX1008, FIG. 6.

Figure 25 likewise shows captured images with annotations for road, ego-lane, and lane instance. EX1008, 24.



EX1008, FIG. 25.

Westmacot also identifies the horizon line as an overlaid feature: "the plane … perpendicular to this road surface normal vector … will provide a line … that would match the horizon." EX1008, 46, FIG. 16. It further teaches projecting the 3D road model into the image plane so "those images can be efficiently annotated … [and] … overlaid on the displayed image." EX1008, 73.

Accordingly, Westmacot discloses overlaying horizon and lane lines on roadway images to generate overlaid images. EX1003, ¶337.

5. [1D] "transmitting the overlaid image to a computing device over a network"

Westmacot in view of Tal renders this element obvious. EX1003, ¶¶338-346.

a. Westmacot

Westmacot discloses creating and storing roadway images with overlaid annotations, but does not expressly transmit them over a network. Westmacot teaches that "the processor 304 receives the captured images from the image capture device 302, and stores them in a memory 306, from which they can be retrieved for use," and that "the annotation data may be displayed with the image (e.g. overlaid on the image)" EX1008, 23, FIGs. 3, 6, 25. Its rendering component "can render … captured images … overlaid with … annotation data A(n)," and annotation data is stored in electronic storage 414 "from which it can be accessed or retrieved." EX1008, 35. Thus, annotated images exist digitally within Westmacot's system. EX1003, ¶339.

b. Tal

Tal complements Westmacot by teaching network transmission of annotated images through a bidirectional vehicle-vision pipeline. Vehicle-mounted cameras capture "digital images 16," and either "unprocessed image(s) 16" or "resultant processed image data 20" are sent to a server in real time. EX1009 ¶[0027]; EX1003 ¶341. Tal emphasizes two-way communication: processed image data and

sensor data "is transmitted to the server 107 as object data 21," while unprocessed images may also be sent for server-side processing. EX1009 ¶¶[0099], [0101]; EX1003 ¶341.

Server processing performs neural inference to produce "bounding boxes 1002, polygons 1003, masks 1004, and landmarks 1005" as part of the processed image data, i.e., annotated imagery. EX1009 ¶[0114]. The server then "communicate[s] with Client(s) 1208 … allow[ing] users to handle the incidents 12 using a user interface 1209," displaying results in map, list, or gallery views. EX1009 ¶¶[0138], [0148]–[0153].

In sum, Tal discloses a feedback loop where raw or partially processed images are uploaded to the server, further processed, and the annotated results returned to client devices (including the vehicle). A POSA would recognize this as a conventional mechanism for transmitting overlaid images and related data in autonomous-driving systems. EX1003 ¶¶342-343.

c. *KSR*

A POSA would have known *how* and *why* to combine these complementary teachings. Section VII.A; EX1003, ¶344. Westmacot generates and stores annotated images. Tal teaches transmitting such images across wireless/cellular networks. Network-based sharing of annotated images was standard in autonomous driving pipelines for training, validation, and review. Applying Tal's routine

transmission mechanisms to Westmacot's digital overlays would have been the predictable use of known techniques, with a reasonable expectation of success. Section VII.A; EX1003, ¶¶345-346.

Accordingly, Westmacot in view of Tal renders obvious "transmitting the overlaid image to a computing device over a network."

6. [1E] "receiving a modification of the line from the computing device, the modification comprising a new line at a second position"

Westmacot in view of Tal renders this element obvious. EX1003, ¶¶347-358.

a. Westmacot

Westmacot teaches user-driven modification of overlaid lines via its annotation interface (UI 412): "the user can modify the annotation data A(n) … to better align the assumed lane boundaries R1, R2 with the actual lane boundaries." EX1008, 33, 37, FIGs. 6, 18–25. Such edits generate "a new line at a second position" that propagates through the dataset. Westmacot also overlays the horizon line: "the plane … perpendicular to this road surface normal vector … provide[s] a line … that would match the horizon." EX1008, 46, FIG. 16. A POSA would have recognized the same modification mechanism applies to horizon lines as to lane lines. EX1003, ¶350. Thus, Westmacot supplies the "what": user-generated line

modifications that create a new line position, but does not expressly disclose the network-based receipt of the user's line modification.

b. Tal

Tal supplies the how: "over a network." Tal reinforces Westmacot by teaching a bidirectional vision pipeline where vehicle cameras capture "digital images 16" and transmit either "unprocessed image(s) 16" or "processed image data 20" to a server. EX1009 ¶[0027]. The server receives "processed image data 20 and sensor data 17 … as object data 21," performs inference to generate "bounding boxes 1002, polygons 1003, masks 1004, [and] landmarks 1005," and communicates results to client devices via a user interface. EX1009 ¶¶[0099], [0101], [0114], [0138], [0148]–[0153]. In short, Tal discloses a feedback loop where raw or processed images are sent to the server and annotated results returned, i.e., conventional mechanisms for transmitting overlaid images in AV systems. EX1003 ¶351.

c. *KSR*

Sending Westmacot's edited line through Tal's channel so the peer node receives that modification is a predictable combination of complementary teachings. It is simply applying a known technique in an analogous context: Westmacot's edits (endpoints, deltas, spline coefficients, or a replace-line message) are routine payloads Tal already carries. A POSA would be motivated by

scalability (server offload), collaboration/consistency (shared remote annotation), and seamless cloud integration, and would reasonably expect success, yielding exactly the claimed receipt of a modified line at a second position. Section VII.A.

Accordingly, the combination renders obvious "a new line at a second position." EX1003, ¶¶352-358.

7. [1F] computing a camera parameter based on the new line

Westmacot in view of Tal renders this element obvious. EX1003, ¶¶359-367.

a. Westmacot

Westmacot discloses computing camera parameters from the geometry of detected or modified lines. Westmacot's "parameter computation component 408" calculates values such as camera height (H), lateral offset (S), forward point, horizon line, and lane width (W), all derived from roadway line geometry. EX1008, 34–37, 42–49, FIG. 3. These parameters update when lane boundaries are adjusted to keep the 3D road model consistent with captured images, for example, adjusting lane width W adapts the road model and automatically propagates new annotations. EX1008, 37. Westmacot also explains that camera orientation is derived from lane tangents, the horizon line from intersecting the road-surface normal with the image plane, and camera height/offset from reconstructed road geometry. EX1008, 29–30, 34–35, 39–40, 47, FIGs. 15–16. Recalculating camera

parameters obviously follows directly from line modifications. EX1003, ¶¶360-361. But Westmacot does not expressly disclose distributing and recomputing calibration data across a network.

b. Tal

Tal expressly teaches distribution, so even if PO argues Westmacot computes calibration locally, Tal teaches performing those computations remotely and sharing results across the network. Specifically, Tal discloses transmitting roadway imagery and detection/overlay data from the in-vehicle device to a remote server, either as "unprocessed image(s) 16" or "processed image data 20", over wireless or cellular links. EX1009, ¶¶[0027], [0030], [0052], [0075]). At the remote site, the server processes the received images to produce calibration-related outputs or modified image information, which can be accessed via a web application for annotation, calibration, or review. EX1009, Summary, FIG. 12, ¶[0138]. These teachings confirm a distributed calibration framework in which remote operations complement Westmacot's local computations and make results available back to the vehicle or use. EX1003, ¶362.

c. *KSR*

A POSA would have known *how* and *why* combining these disclosures is a routine application of known distributed processing. The motivations are clear: (i) verification and consistency across systems, (ii) reduced onboard processing load,

and (iii) enabling human-in-the-loop calibration with remote updates returned in real time. EX1003, ¶¶363-367.

Accordingly, Westmacot discloses computing camera parameters from line geometry, and in view of Tal, it would have been obvious to implement this in a networked architecture where recalculations occur locally or remotely, the combination being nothing more than the predictable use of known techniques yielding updated parameters from new line positions. Section VII.A; *KSR*, 550 U.S. at 416; *Wyers*, 616 F.3d 1231, 1240 (Fed. Cir. 2010) ("[O]bviousness may be found where a combination of familiar elements according to known methods yields no more than predictable results."); EX1003, ¶367.

8. [1G] transmitting data representing the camera parameter to the camera device.

Westmacot in view of Tal renders this element obvious. EX1003, ¶368.

a. Westmacot

Westmacot computes camera parameters (height H, offset S, lane width W, forward point, horizon line, orientation) using its "parameter computation component 408" and shares them with model adaptation 410, rendering 416, and UI 412 to update the 3D road model. EX1008, 34–37, 42–49, FIGs. 3, 6, 25, 30–32. PO may argue Westmacot does not expressly disclose transmitting parameters back to the camera device (302). Tal fills this gap. EX1003, ¶369.

b. Tal

Tal teaches a two-way, networked pipeline for image-derived calibration data: "digital images 16 … and … resultant processed image data 20 … [are] transmitted to the server 107," EX1009, ¶[0027]; processing may be on-device, on the server, or shared "in FIG. 4 … shared/distributed processing environment 400," with "the server 107 … organizing, storing, processing and disseminating … object data 21," EX1009, ¶¶[0030], [0138]. It further clarifies that "some or all of image adjustment may be performed on the vehicle" and a "second portion … on the server 107," yielding "bandwidth transmission savings" when only refined outputs are returned. EX1009, ¶¶[0052], [0075]. Thus, Tal supplies the missing implementation for computing calibration remotely and returning results to the originating camera device. EX1003, ¶370.

c. *KSR*

A POSA would have known *how* and *why* to extend Westmacot's internal parameter sharing to the camera device using Tal's feedback loop. The benefits are clear: enabling self-calibration. EX1009, ¶[0063]), reducing onboard load through distributed processing, and ensuring consistent results. EX1003, ¶¶371-374. This is a predictable use of known techniques with a reasonable expectation of success. Section VII.A; EX1003, ¶375.

C. Claim 2

Claim 2 recites that the "camera parameter comprises one of camera height, viewing angle, and road plane normal." Westmacot discloses all three. EX1003 ¶¶376-377.

Camera height: Westmacot computes "the height H of the camera 302 above the road," either measured directly or derived from a reconstructed surface mesh averaging the camera path's height. EX1008, 28, 32, 34, 47, FIGs. 3, 3A, 5; EX1003 ¶378.

Viewing angle: The system captures camera orientation as a "forward point" and "horizon line," where the forward point is "the pixel the car appears to head toward" and the horizon line derives from the road-surface normal projected onto the image plane. EX1008, 35, 42–43, 46; EX1003 ¶379.

Road plane normal: Westmacot estimates the road normal n by taking the cross product of motion vectors in the road plane. EX1008, 28, FIG. 27; EX1003 ¶380.

Thus, Westmacot calculates camera height, viewing angle, and road plane normal as part of its calibration process, satisfying claim 2. EX1003 ¶381.

D. Claim 3

Claim 3 recites "receiving an image of the roadway while the vehicle is moving." Westmacot discloses this. EX1003 ¶382.

Westmacot teaches "receiving a time sequence of two-dimensional images … captured by an image capture device of [a] travelling vehicle" to "reconstruct … a path travelled by the vehicle." EX1008, 5, 9. Figure 3 shows "vehicle 300 … used to capture road images … frames of short video segments recorded as the vehicle drives along a road." EX1008, 24. The system relies on imagery captured in motion, "driving the car is itself a form of annotation," and uses dashcam videos filtered to include "only frames … at least 1m apart according to GPS." EX1008, 20–21, 24.

Thus, Westmacot teaches receiving roadway images while the vehicle is moving, satisfying claim 3. EX1003 ¶383.

E. Claim 4

Claim 4 recites that "receiving the image of the roadway while the vehicle is moving comprises detecting that the vehicle is traveling above a pre-defined speed or for a pre-defined duration." Westmacot in view of Tal renders claim 4 obvious. EX1003, ¶384.

Westmacot filters images so frames are captured only when the vehicle is moving: "to remove parts where the car moves very slow or stands still … only

frames that are at least 1m apart according [to] GPS are included." EX1008, 24. A POSA would recognize this distance cutoff as equivalent to applying a minimum speed or duration threshold, since fixed-interval sampling excludes idle or low-speed periods. EX1003 ¶385.

It would have been obvious to use speed or time thresholds instead of distance, as autonomous driving systems routinely gated imagery using telemetry (distance, speed, or time) to exclude non-driving frames. EX1003 ¶¶386–387.

Even if not express in Westmacot, Tal confirms this by transmitting roadway images with telemetry including position, speed, and motion data. EX1009 ¶¶[0027], [0030], [0052], [0075], [0138].

A POSA would have known *how* and *why* to combine Westmacot's distance-based filtering with Tal's speed/time metadata to ensure images reflect meaningful driving conditions, thereby improving calibration accuracy, training consistency, and system efficiency. EX1003 ¶¶389–392.

F. Claim 5

Claim 5 recites that transmitting the overlaid image comprises transmitting it to a web-based application. Westmacot in view of Tal renders this claim obvious. EX1003, ¶393.

Westmacot discloses capturing roadway images, overlaying annotation data, and transmitting that data within its system for training and calibration. EX1008

FIG. 3 (UI 412; rendering 416; storage 414). While it produces digitized overlaid images suitable for transmission, it does not expressly explain how end users access them. EX1003 ¶¶394–395.

Tal fills this potential gap by teaching that transmitted images are accessible via a web interface: servers "serve as gateway to users via web access and present data … in a meaningful and intuitive manner," and users may log in through "a user interface … either a web application … or a client/server application." EX1009 Summary, ¶[0138], FIG. 12; EX1003 ¶¶396-397.

A POSA would have known *how* and *why* to combine. EX1003 ¶¶398–402. Westmacot provides annotated images; Tal offers the conventional web-based mechanism for distributing and viewing them. By the priority date, web delivery was standard for remote collaboration, offering scalability, platform independence, and low deployment cost. Integrating Tal's web interface into Westmacot's image pipeline would predictably enhance accessibility, support distributed annotation, and streamline training.

Thus, applying Tal's web-access model to Westmacot's system represents a routine, predictable improvement with a reasonable expectation of success, rendering claim 5 obvious. EX1003 ¶403.

G. Claim 6

Claim 6 recites receiving "an identification of a new horizon line from the computing device." Westmacot in view of Tal renders this obvious. EX1003, ¶404.

Westmacot teaches automatic horizon-line computation by projecting the road-surface normal onto the image plane. EX1008, 46, FIG. 16; EX1003, ¶¶405-406. It notes that parameters may require manual correction during the "human fixer" stage. EX1008, 49. Its annotation interface (UI 412) allows users to adjust roadway feature lines such as lane boundaries, with edits propagated through the 3D road model. EX1008, 35, 51, FIGs. 6, 25. A POSA would recognize that the same mechanism applies to correcting the horizon line—user adjustments in the interface would similarly update the overlay and trigger parameter recalculation to maintain consistent camera calibration. EX1003 ¶407.

Even if not express, Tal reinforces this by teaching that roadway images and calibration parameters are presented through computing-device interfaces for user input, e.g., "device 101 … [with] user interface 119 … display … orientation … settings, parameters." EX1009, UI 119, ¶[0138], ¶[0035]; EX1003, ¶408. Tal describes manual roadway-feature identification on tablets and smartphones, confirming such interactions were conventional.

A POSA would have known *how* and *why* to combine, extending Westmacot's interface from lane-line edits to horizon-line adjustment. Both are

geometric calibration primitives, and enabling horizon-line correction would predictably reduce pitch misalignment and improve accuracy. EX1003 ¶¶409–414.

Accordingly, Westmacot in view of Tal renders claim 6 obvious as a straightforward application of known techniques with a reasonable expectation of success. EX1003 ¶414.

H. Claim 7

Westmacot in view of Tal renders claim 7 obvious. EX1003, ¶415.

Westmacot's component 408 computes and updates camera parameters (height H, offset S, orientation, lane width W) when line geometry is adjusted—e.g., "a user can adjust … W … [and] the annotation data … [is] automatically adapted." EX1008, 34–35, 49, FIG. 4; EX1003 ¶¶416-417.

Even if not express, Tal teaches a bidirectional pipeline where raw or processed imagery and calibration data are exchanged between vehicles and servers. EX1009 ¶¶[0027], [0030], [0052], [0075], [0138], FIG. 12; EX1003, ¶418.

A POSA would have combined these teachings. EX1003 ¶419. Westmacot shows parameter recalculation following line edits, while Tal shows bidirectional data transmission. Together, they suggest transmitting recalculated parameters back to the vehicle, completing the calibration loop. The motivations are clear: (i) calibration feedback: updating camera settings or embedding metadata (EX1009 ¶[0063]); (ii) consistency: servers distributing recalculated parameters across

devices (EX1009 ¶¶[0027], [0030], [0052], [0075]); and (iii) predictability: feedback loops returning calibration data were conventional by the priority date. EX1003 ¶¶419–423.

Accordingly, combining Westmacot and Tal renders claim 7 obvious as a routine application of known elements with a reasonable expectation of success. EX1003 ¶¶419–423; Section VII.A.

I. Claims 8-20 are obvious over Westmacot and Tal.

Claims 8-20 are obvious over Westmacot and Tal. EX1003, ¶¶424-425. These claims merely restate the substantive limitations of claims 1-7 in alternative statutory forms, i.e., as computer-readable medium and device claims. Accordingly, because claims 8-20 add no substantive limitations beyond those already rendered obvious by Westmacot in view Tal for claims 1-7, they too are obvious. EX1003, ¶425.

IX. NO OBJECTIVE INDICIA OF PATENTABILITY

Petitioner is unaware of any objective indicia of non-obviousness. If PO alleges objective indicia in its POPR, Petitioner requests a response.

X. MANDATORY NOTICES

A. Real Party in Interest

Petitioner Samsara Inc. ("Samsara") is the only real party-in-interest.

B. Related Matters

Motive asserted the '276 patent in the Northern District of California (Case No. 3:24-cv-00902) in an amended complaint filed on July 9, 2025. EX1013. In the same amended complaint, Motive also asserted U.S. Patent No. 12,062,243 (the "'243 Patent"). The '276 patent is a child of the '580 patent. An IPR trial (IPR2025-00574) was instituted against the '580 patent on August 27, 2025.

A. Lead and Backup Counsel

Lead Counsel	Back-up Counsel
Jason D. Eisenberg (Reg. No. 43,447) STERNE, KESSLER, GOLDSTEIN & FOX PLLC jasone-PTAB@sternekessler.com 1101 K Street, NW, 10th Floor Washington, DC 20005 202.371.2600 (reception) 202.371.2540 (facsimile)	Lestin Kenton (Reg. No. 72,314) Tyler Dutton (Reg. No. 75,069) Jean Selep (Reg. No. 83,022) lkenton-PTAB@sternekessler.com tdutton-PTAB@sternekessler.com jselep-PTAB@sternekessler.com 1101 K Street, NW, 10th Floor Washington, DC 20005 202.371.2600 (reception) 202.371.2540 (facsimile)

Please address all correspondence to lead and backup counsel at the email addresses shown above and PTAB@sternekessler.com. Samsara consents to electronic service by email.

XI. GROUNDS FOR STANDING

Samsara certifies the '276 patent is available for IPR and that Samsara is not barred or estopped from requesting IPR of the '276 patent challenging the patent claims on the grounds identified in this Petition.

XII. CONCLUSION

Grounds 1, 2, 3A/B, and 4 show that the '276 patent inventors did not originate calibrating vehicle-mounted cameras using horizon lines, lane lines, or human-in-the-loop annotation, nor were they first to disclose recalculating camera parameters when those lines are modified. Such calibration methods, including automatic horizon and lane detection, operator correction of overlays, and parameter updates, were well known long before the '276 patent's priority date. The Board should therefore institute inter partes review and cancel the challenged claims.

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg, Reg. No. 43,447
Lestin Kenton, Reg. No. 72,314
Tyler Dutton, Reg. No. 75,069
Jean Selep, Reg. No. 83,022
Counsel for Petitioner Samsara Inc.

Date: October 17, 2025

1101 K Street NW
10th Floor
Washington, DC 20005
(202) 371-2600

CERTIFICATE OF COMPLIANCE WITH TYPE-VOLUME LIMITATION, TYPEFACE REQUIREMENTS, AND TYPE STYLE REQUIREMENTS

1. This Petition complies with the type-volume limitation of 14,000 words, comprising 13,930 words, excluding the parts exempted by 37 C.F.R. § 42.24(a).

2. This Petition complies with the general format requirements of 37 C.F.R. § 42.6(a) and has been prepared using Microsoft® Word in 14-point Times New Roman.

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg
Registration No. 43,447
Counsel for Petitioner Samsara Inc.

Date: October 17, 2025

1101 K Street NW
10th Floor
Washington, DC 20005
(202) 371-2600

CERTIFICATION OF SERVICE (37 C.F.R. §§ 42.6(e), 42.105(a))

The undersigned hereby certifies that the above-captioned **PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 12,136,276** and all supporting exhibits were served via FedEx Express® on October 17, 2025, in their entireties on the following:

GREENBERG TRAURIG (NY)
One Vanderbilt Avenue
New York, NY 10017
Correspondence Address for U.S. Patent No. 12,136,276

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg
Registration No. 43,447
Counsel for Petitioner Samsara Inc.

Date: October 17, 2025

1101 K Street NW
10th Floor
Washington, DC 20005
(202) 371-2600

EXHIBIT B

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSARA INC.,
Petitioner

v.

MOTIVE TECHNOLOGIES, INC.,
Patent Owner

Case IPR2026-00108
U.S. Patent No. 12,062,243

PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 12,062,243

Mail Stop "PATENT BOARD"
Patent Trial and Appeal Board
U.S. Patent & Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

I. INTRODUCTION 1
II. THE '243 PATENT 1
A. Overview 1
B. Level of Ordinary Skill in the Art 2
C. Claim Construction 2
III. IDENTIFICATION OF THE CHALLENGE (37 C.F.R. § 42.104(b)) 3
A. Prior Art 3
B. Grounds 4
IV. OVERVIEW OF THE PRIOR ART 4
A. Chun 4
B. Chao 5
C. Shanmugamani 6
D. Zhengyang 6
E. He 7
V. GROUND 1: CHUN IN VIEW OF CHAO RENDERS OBVIOUS CLAIMS 1, 2, 6-14, 18-20, AND 24-33. 8
A. Independent Claim 1 9
1. [1Pre] 9
2. [1A] 9
3. [1B] and [1C] 10
4. [1D] 12
a. Chun 12
b. Chao 12
c. Motivation to Combine (MTC) 13
5. [1E] 15
6. [1F] 16
7. [1G] 17
a. Chun 17

b. Chao 18
c. MTC 19
8. [1H] 21
a. Chun 21
b. Chao 22
c. MTC 22
B. Claim 2 23
1. Chun 23
2. Chao 24
3. MTC 24
C. Claim 6 25
1. Chun 25
2. Chao 26
3. MTC 26
D. Claims 7 and 8 27
1. Chun 27
2. Chao 28
3. MTC 28
E. Claim 9 29
F. Claim 10 31
G. Claim 11 31
H. Claim 12 32
I. Claims 13, 14, 18-20, 24, and 25 33
J. Independent Claim 26 34
K. Claim 27 35
L. Claim 28 35
M. Claim 29 36
N. Claims 30-33 36

VI. GROUND 2: CHUN IN VIEW OF CHAO AND FURTHER IN VIEW OF SHANMUGAMANI RENDERS OBVIOUS CLAIMS 3-5, 10, 15-17, AND 21-23. 37
A. Claim 3 38
1. Chun-Chao 38
2. Shanmugamani 39
3. MTC 41
B. Claim 4 42
1. Chun-Chao 42
2. Shanmugamani 42
3. MTC 46
C. Claim 5 47
1. Chun-Chao 47
2. Shanmugamani 47
3. MTC 51
D. Claims 15-17 and 21-23 51
VII. GROUND 3: ZHENGYANG IN VIEW OF HE RENDERS OBVIOUS CLAIMS 1-6, 9-18, AND 20-24. 53
A. Independent Claim 1 55
1. [1Pre] 55
2. [1A] 55
3. [1B] and [1C] 55
4. [1D] 57
5. [1E] 58
6. [1F] 59
a. Zhengyang 59
b. He 59
c. MTC 60
7. [1G] 61
a. Zhengyang 61

b. He and MTC 62
8. 1[H] 62
B. Claim 2 63
C. Claim 3 65
D. Claim 4 66
E. Claim 5 68
1. Zhengyang 68
2. He 68
F. Claim 6 70
G. Claim 9 72
1. Zhengyang 72
2. He 72
3. MTC 73
H. Claim 10 73
I. Claim 11 74
J. Claim 12 75
K. Claims 13-18 and 20-24 76
VIII. GROUND 4: ZHENGYANG IN VIEW OF HE AND FURTHER IN VIEW OF CHAO RENDERS OBVIOUS CLAIMS 7-8, 19, 25-33. 76
A. Claims 7, 8, 19 and 25 76
1. Zhengyang-He 77
2. Chao 78
3. MTC 79
B. Claims 26-33 80
IX. NO OBJECTIVE INDICIA OF PATENTABILITY 81
X. MANDATORY NOTICES 82
A. Real Party in Interest 82
B. Related Matters 82
A. Lead and Backup Counsel 82

XI. GROUNDS FOR STANDING..... 83
XII. CONCLUSION..... 83

PETITIONER'S EXHIBIT LIST

Exhibit	Description
1001	U.S. Patent No. 12,062,243 to Hassan *et al.* ("Hassan")
1002	Prosecution History of U.S. Patent No. 12,062,243
1003	Declaration of Dr. Trevor Darrell
1004	*Curriculum Vitae* of Dr. Trevor Darrell
1005	Chun, Sehyun, et al. "NADS-Net: A nimble architecture for driver and seat belt detection via convolutional neural networks." Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. 2019.
1006	U.S. Publication No. 2020/0242381 to Chao *et al.* ("Chao")
1007	Chinese Publication No. CN111860253 to Dongguan Zhengyang Electronic Mechanical Co ltd ("DZE")
1008	He, K., Gkioxari, G., Dollar, P., and Girschick, R. "Mask R-CNN" IEEE International Conference on Computer Vision (2017)
1009	Caruana, Rich. "Multitask Learning." Machine Learning 28 (1997): 41-75.
1010	Shanmugamani, R. Deep Learning for Computer Vision (2018)
1011	P. Smith, M. Shah and N. da Vitoria Lobo, "Determining driver visual attention with one camera," in IEEE Transactions on Intelligent Transportation Systems, vol. 4, no. 4, pp. 205-218, Dec. 2003, doi: 10.1109/TITS.2003.821342.
1012	Baheti, Bhakti & Gajre, Suhas & Talbar, Sanjay. (2018). Detection of Distracted Driver Using Convolutional Neural Network. 1145-11456. 10.1109/CVPRW.2018.00150.
1013	M. S. Majdi, S. Ram, J. T. Gill and J. J. Rodríguez, "Drive-Net: Convolutional Network for Driver Distraction Detection," 2018 IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI), Las Vegas, NV, USA, 2018, pp. 1-4, doi: 10.1109/SSIAI.2018.8470309.
1014	U.S. Patent No. 11,176,384 B1 to Yang et al. ("Yang")
1015	U.S. Patent No. 12,347,177 B2 to Li et al. ("Li")

Exhibit	Description
1016	R. Ranjan, V. M. Patel and R. Chellappa, "HyperFace: A Deep Multi-Task Learning Framework for Face Detection, Landmark Localization, Pose Estimation, and Gender Recognition," in IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 41, no. 1, pp. 121-135, 1 Jan. 2019, doi: 10.1109/TPAMI.2017.2781233
1017	Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 580-587).
1018	Girshick, R. (2015). Fast R-CNN. In Proceedings of the IEEE international conference on computer vision (pp. 1440-1448).
1019	Ren, S., He, K., Girshick, R., & Sun, J. (2016). Faster R-CNN: Towards real-time object detection with region proposal networks. IEEE transactions on pattern analysis and machine intelligence, 39(6), 1137-1149.
1020	Li, M., Chen, S., Chen, X., Zhang, Y., Wang, Y., & Tian, Q. (2019). Actional-structural graph convolutional networks for skeleton-based action recognition. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 3595-3603).
1021	Christoph Maurer, Integration and Deployment of Machine Learning Models (Master's Thesis, Graz University of Technology, Feb. 2021)
1022	Umar Zakir Abdul, Hamid; et al. (2016). "Current Collision Mitigation Technologies for Advanced Driver Assistance Systems-A Survey". PERINTIS eJournal. 6 (2).
1023	Bojarski, M., Yeres, P., Choromanska, A., Choromanski, K., Firner, B., Jackel, L., & Muller, U. (2017). Explaining how a deep neural network trained with end-to-end learning steers a car. arXiv preprint arXiv:1704.07911.
1024	G. Gkioxari, B. Hariharan, R. B. Girshick, and J. Malik. R-cnns for pose estimation and action detection. CoRR, abs/1406.5212, 2014
1025	LeCun, Y., Bengio, Y., & Hinton, G. (2015). "Deep learning." Nature, Vol. 521 (pp. 436-444).
1026	Ioffe, S., & Szegedy, C. (2015). Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift.

Exhibit	Description
1027	Goodfellow, I., Bengio, Y. & Courville, A. (2016). Deep Learning.
1028	He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving Deeping into Rectifiers: Surpassingb Human-Level Performance on ImageNet Classification.
1029	Hu, J., Shen, L., & Sun, G. (2018). Squeeze-and-Excitation Networks. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 7132-7141).
1030	He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition.
1031	Krizhevsky, A., Sutskever, I., & Hinton, G. (2012). ImageNet Classification with Deep Convolutional Neural Networks.
1032	Abouelnaga, Y., Eraqi, H., & Moustafa, M. (2017). Real-time Distracted Driver Posture Classification. In 32nd Conference on Neural Information Processing Systems (NIPS 2018), Montreal, Canada.
1033	Kose, N., Kopuklu, O., Unnervik, A., & Rigoll, G. (2019). Real-Time Driver State Monitoring Using a CNN Based Spatio-Temporal Approach.
1034	Braunagel, C., Stolzmann, W., Kasneci, E., & Rosensteil, W. (2017). Driver Activity Recognition in the Context of Conditionally Autonomous Driving. In 2015 IEEE 18th International Conference on Intelligent Transportation Systems (pp. 1652-1657).
1035	Engström, J., et al. 2019. A conceptual framework and taxonomy for understanding and categorizing driver inattention.
1036	Choi, I., & Kim, Y. (2015). Head Pose and Gaze Direction Tracking for Detecting a Drowsy Driver. In Applied Mathematics & Information Sciences 9, No. 2L (pp. 505-512).
1037	*Motive Techs., Inc. v. Samsara Inc.*, Case No. 24-cv-00902-JD (NDCA), Second Amended Complaint
1038	Simonyan, K. & Zisserman A. 2015. Very Deep Convolutional Networks For Large-Scale Image Recognition.
1039	Maurer, C., Integration and Deployment of Machine Learning Models. Graz University of Technology Master's Thesis. (2021).

Exhibit	Description
1040	Boureau et al. 2010. A Theoretical Analysis of Feature Pooling in Visual Recognition, Proceedings of the 27th International Conference on Machine Learning.
1041	Declaration of Dr. Sylvia Ellis-Hall

I. INTRODUCTION

U.S. Patent No. 12,062,243 offers no technical advance. It adopts the same backbone-plus-heads CNN (convolutional neural network) template long used in ADAS (advanced driver assistance system) driver-monitoring: a shared backbone, multiple task heads, supervised loss minimization, and storing a trained model for inference. Ground 1 (Chun-Chao) shows this architecture and storage were conventional. Ground 2 (Chun-Chao-Shanmugamani) shows the claimed detection and pose heads were routine. Ground 3 (Zhengyang-He) confirms modular heads were predictably added to a shared backbone. Ground 4 (Zhengyang-He-Chao) shows no novelty in action-based outputs. The claims redeploy known CNN components warranting institution.

II. THE '243 PATENT

A. Overview

The '243 patent describes a well-known, multitask training pipeline. It frames distracted-driver detection not as a simple classifier, but as a multitask, computer-implemented architecture that trains a backbone CNN on labeled in-cabin driver images and then deploys the trained model in-vehicle for real-time distraction classification. It broadens "distraction" beyond mobile-phone use to include drinking, smoking, talking to passengers, interacting with controls, leaning out of the cabin, and other secondary behaviors, and the claims recite learning those signals jointly. It uses a shared backbone feature extractor coupled to

multiple task-specific prediction heads, including a distraction-classification head, object-detection head, and pose-estimation head. During training, labeled driver images are input into the network, each head computes a loss, and those head losses flow through the backbone to update backbone parameters. After convergence, the backbone and head parameters are stored as the "distraction classification model" and used to classify distraction states in real time on a dash-mounted device. EX1003, ¶¶23-39

B. Level of Ordinary Skill in the Art

A person of ordinary skill in the art ("POSITA") at the time of the alleged invention would have had a bachelor's in computer science with a focus on computer/machine vision, and at least two years of experience with each of: (1) machine-learning (including neural-network) methods as applied to machine/computer vision and (2) classical machine/computer vision algorithms (e.g., R-CNN, Mask R-CNN, Fast R-CNN, and SSD). EX1003, ¶86. Additional education could serve as a substitute for experience. *Id.* The Petition is supported by the Declaration of Dr. Trevor Darrell. EX 1003, ¶1-22

C. Claim Construction

Petitioner construed the terms according to their ordinary and customary meanings. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312-13 (Fed. Cir. 2005).

III. IDENTIFICATION OF THE CHALLENGE (37 C.F.R. § 42.104(b))

A. Prior Art

The Grounds rely on the following prior art, not applied during prosecution. EX1003, ¶¶75-85:

Reference	Exhibit(s)	Date	Prior Art
Chun	EX1005	Published Feb. 2, 2020	§102(a)(1)
Chao	EX1006	Filed March 26, 2020; Published July 30, 2020	§102(a)(1)-(2)
Zhengyang	EX1007	Filed July 10, 2020; Published October 30, 2020	§102(a)(1)-(2)
He	EX1008	Published June 5, 2018	§102(a)(1)
Shanmugamani	EX1010	Published February 23, 2018	§102(a)(1)

Petitioner has submitted a declaration from Dr. Hall-Ellis—a librarian with over forty years of professional cataloging and bibliographic-verification experience—to confirm the public accessibility and authenticity of Chun, He (Mask R-CNN) and Shanmugamani. *See* EX1041. As explained in her declaration, the Chun paper was published in the Proceedings: 2019 International Conference on Computer Vision and was publicly available no later than February 2, 2020. EX1041, ¶¶53-59. He was published in the IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 42, Issue 2, and was publicly available no later than June 5, 2018, consistent with the date reflected on IEEE Xplore. EX1041, ¶¶46-52. Shanmugamani was issued by Packt in 2018 and was publicly available

no later than February 23, 2018. EX1041, ¶¶67-72. Petitioner relies on Zhengyang, originally published in Chinese. A complete copy of the original Zhengyang document, together with a certified English translation and the translator's certificate of accuracy, is provided in a single combined Exhibit 1007.

B. Grounds

Ground	Basis	Reference(s)	Claims
1	§103	Chun and Chao	1, 2, 6-9, 11-14, 18-20, and 24-33
2	§103	Chun, Chao, and Shanmugamani	3-5, 10, 15-17, and 21-23
3	§103	Zhengyang and He	1-6, 9-18, 20-24
4	§103	Zhengyang, He and Chao	7-8 19, 25-33

IV. OVERVIEW OF THE PRIOR ART

Detecting distracted driving using a multi-task CNN with a backbone and prediction heads for object detection, pose estimation, and distraction classification was well known. EX1003, ¶40.

A. Chun

Chun teaches (National Advance Driver Simulator) NADS-Net, a lightweight CNN for in-vehicle monitoring because "most accidents are caused by distracted driving." EX1005, 1. Subjects were recorded in-cabin (dual cameras&IR) performing distraction behaviors such as "drinking, using smart phones, yawning," and the frames were manually annotated (seatbelt masks&MS-COCO keypoints). EX1005, 3-4, FIGs. 1-2. NADS-Net uses a ResNet-50 FPN

backbone with multiple heads (keypoint, PAF, seatbelt). EX1005, 5, FIG. 3.

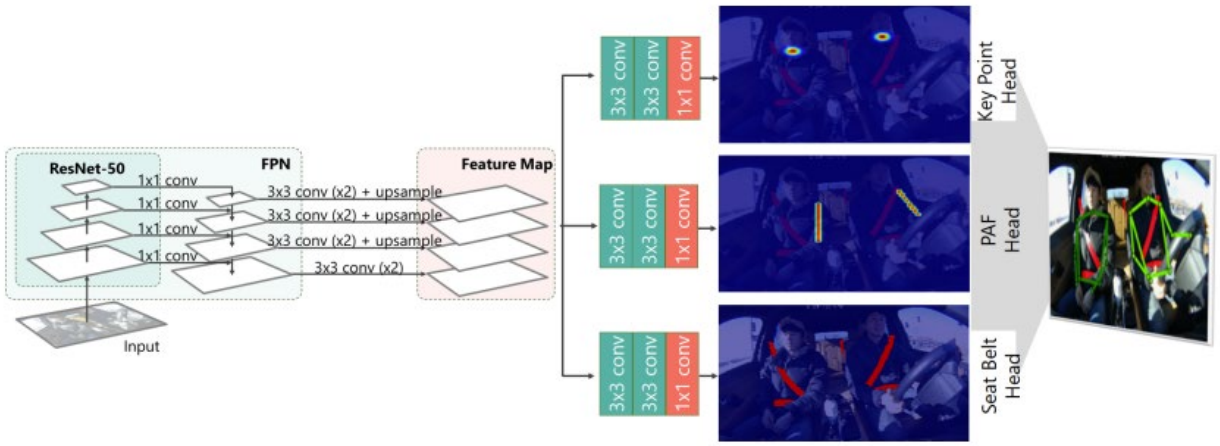


Figure 3: NADS-Net architecture.

EX1005, FIG. 3.

Each head uses "two 3×3 convolutions and a 1×1 convolution to predict a pixel-wise probability." EX1005, 5. The trained model is then used for IVMS (in vehicle monitoring system), which "play a critical role" in mitigating accidents. EX1005, 1-2. This is standard supervised CNN training (joint loss → backbone updates) and storage for inference. EX1003, ¶¶41-48.

B. Chao

Chao teaches monitoring occupants and predicting categorical distraction states (e.g., "Drink," "Call," "Text"). EX1006, ¶¶[0098], [0101], [0113]. Features are extracted by a 2D CNN backbone and fed to a spatio-temporal (ST) classifier. Chao teaches supervised training: ST outputs are used as "positive or negative feedback," and the CNN is "configured to be further trained based on the

feedback." EX1006, ¶¶[0098], [0041], [0043], [0046]. After training, weights are stored in memory. EX1006, ¶[0054]. This is a conventional backbone+head CNN trained end-to-end and stored for inference. EX1003, ¶¶49-53.

C. Shanmugamani

Shanmugamani teaches how to build deep-learning CV (cross-validation) models for "classification, detection, retrieval, segmentation, generation, captioning, and video classification," including autonomous driving. EX1010, 1-5. A neural network is "a collection of perceptrons and activation functions," and CNNs use convolution, pooling, batch-norm, activation, FC (full-connected) layers, and backpropagation to "detect features or patterns in images." EX1010, 7, 18. For object detection, Shanmugamani describes R-CNN (region-based CNN), Fast/Faster R-CNN, SSD (single-shot multibox detector), YOLO (you only look once), and TensorFlow implementation with pre-trained COCO (common objects in context) models and retraining. EX1010, 109-123, 123-126. Shanmugamani also teaches pose estimation as localization of body joints. EX1010, 110-119; EX1003, ¶¶54-64.

D. Zhengyang

Zhengyang teaches a unified, end-to-end multi-task CNN for driver-attribute recognition. Prior approaches used multiple separate models; Zhengyang instead "adopt[s] a multi-task network model to simultaneously recognize multiple

features." EX1007, ¶¶[0003], [0005]. Labeled driver images are "input into the multi-task main network to obtain a feature map… [then] input into respective attribute recognition modules." EX1010, ¶¶[0020]-[0025]. RGB (red-green-blue) frames feed a backbone; Zhengyang's OSME module creates four feature maps routed to separate heads. EX1010, ¶[0069]. Loss minimization updates parameters until convergence, then weights are stored in a computer-readable medium. EX1010, ¶¶[0022], [0090], [0044]-[0045]; EX1003, ¶¶65-70.

E. He

He teaches a multi-task CNN with a shared ResNet-FPN (feature pyramid network) backbone and multiple heads, extending Faster R-CNN by adding a mask head and a keypoint (pose-estimation) head. EX1008, 7-8. For pose, He states: "Our framework can easily be extended to human pose estimation…[we] adopt Mask R-CNN to predict K masks," implemented as "eight 3×3 512-d conv layers followed by a deconv layer and 2× bilinear upscaling." *Id.* The head outputs keypoint heatmaps, trained via cross-entropy, and uses joint backpropagation so losses from the pose head update both the head and the shared backbone end-to-end. EX1003, ¶¶71-74.

V. GROUND 1: CHUN IN VIEW OF CHAO RENDERS OBVIOUS CLAIMS 1, 2, 6-14, 18-20, AND 24-33.

Chun in view of Chao discloses each limitation of the independent claims (1, 13, 20, 26, 30), and the dependent claims fall for the same reasons. EX1003, ¶88-89, 97. A POSITA would have combined Chun and Chao because both address distracted-driving detection using multi-task CNNs and teach compatible shared-backbone architectures. EX1003, ¶90. Chun provides the backbone and multiple heads, and Chao supplies the express teaching of standard state-classification and loss-minimizing parameter update. EX1003, ¶91. Chao explains the ST model operates over "features extracted from images by the 2D CNN," EX1006, ¶[0101], which are feature maps of the same type produced by Chun's backbone, so routing Chun's shared features into Chao's classification head would require no architectural redesign. Minimizing loss during training was routine and Chao expressly teaches the ST model "provides feedback to further train the 2D CNN," EX1006, ¶¶[0043], [0098], which a POSITA recognizes as standard backprop-based parameter update. Both references seek improved driver-state accuracy in-vehicle and use the same standardized backbone/head paradigm. EX1003, ¶¶92-95. A POSITA therefore would have reasonably expected success, and the combination yields the claimed subject matter as a predictable application of known CNN practice. EX1003, ¶96.

A. Independent Claim 1

1. [1Pre]

Chun discloses the preamble of claim 1. Chun describes training a CNN called NADS Net that "utilizes the feature pyramid network backbone with multiple detection heads to achieve optimal performance for driver/passenger state detection tasks." EX1005, 1. That training for "2D driver/passenger pose estimation and seat belt detection" is a method, as recited in claim 1. EX1003, ¶98.

2. [1A]

Chun discloses [1A]. Chun describes capturing in-vehicle frames and annotating them so "annotators were instructed to segment all visible seat belts with a binary mask and to mark x and y pixel coordinates of body key points" under the MS COCO protocol. EX1005, 3-4, FIG. 1. These annotations constitute labeled examples identifying seat-belt regions and body key points.

Example annotations are shown below:

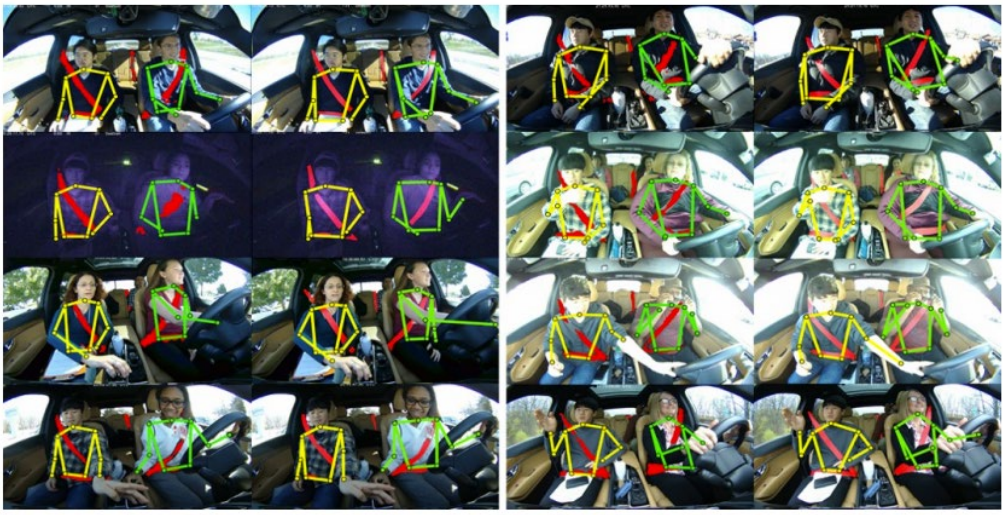


Figure 1: First and third columns in the figure are results produced by NADS-Net and second and fourth columns are the corresponding ground truth annotations.

EX1005, FIG. 1.

Further, the CNN was "pre trained on the MS COCO data set and then transferred to the new driving data set" using Keras with a TensorFlow backend, and trained using labeled key point annotations. EX1005, 5, 7. Chun shows labeled images being input into a neural network for supervised learning, satisfying [1A]. EX1003, ¶¶99-102.

3. [1B] and [1C]

Chun discloses [1B] and [1C]. At the outset, for [1C], a POSITA would understand this limitation is met by a single qualifying head. The phrase "one or more prediction heads" establishes the numerical requirement, and "comprising one or more of" the listed head types indicates any subset satisfies the limitation.

The claim therefore does not require all three head types. A single prediction head of one of the recited types is sufficient. Chun satisfies at least this.

Chun teaches NADS-Net uses a feature-extraction backbone: "a feature pyramid network (FPN) backbone with multiple detection heads to achieve optimal performance for driver/passenger state detection tasks." EX1005, 1. Figure 3 shows the ResNet-50 FPN backbone generating multi-resolution feature maps that are "further convolved with 1×1 convolutions" to form a shared representation. EX1005, 5-6. These shared backbone features feed multiple downstream prediction heads.

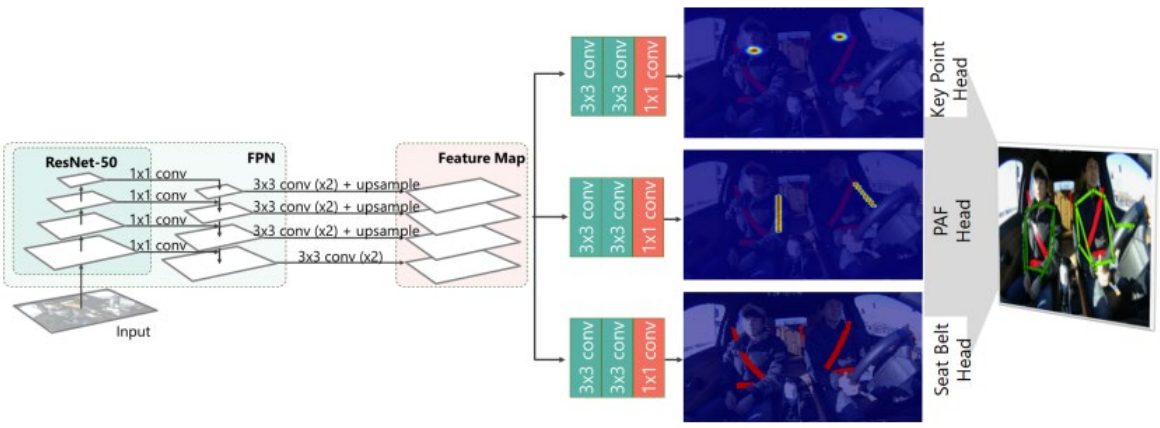


Figure 3: NADS-Net architecture.

EX1005, FIG. 3.

Chun also discloses the prediction heads. NADS-Net "utilizes a FPN backbone with multi-branch detection heads," including key-point, part-affinity field, and seat-belt segmentation heads. EX1005, 1, FIG. 3. Each head "employs two 3×3 convolutions and a 1×1 convolution" to output pixel-wise probabilities:

joints, joint relationships, or seat-belt mask. EX1005, 5.

Chun discloses a shared backbone and multiple prediction heads, satisfying [1B] and [1C]. EX1003, ¶¶103-112.

4. [1D]

Chun-Chao renders obvious [1D].

a. Chun

Chun states "a vast majority of vehicle accidents are caused by distracted driving behaviors" such as "using smartphones, smoking, drinking, eating, and engaging in conversation," and trains NADS-Net on those same behaviors ("drinking, using smart phones… yawning… adjusting sun visor"). EX1005, 1,4. A POSITA would view these labeled behaviors as distraction classes. Though Chun does not name a separate distraction head, adding a soft-max classification branch that consumes the same shared feature maps as the pose and seat-belt heads would have been obvious. The new branch would simply consume the shared feature maps already used by Chun's pose and seat-belt heads and output distraction categories - attentive, phone use, or drinking.

This is also obvious in view of Chao, which expressly discloses a distraction-classification head. EX1003, ¶¶114-119.

b. Chao

Chao describes a CNN extracting spatial features from video frames and a spatio-temporal model classifying driver "states" including "paying attention,

texting, drinking, [and] looking at the radio." EX1006, ¶¶[0094]-[0098]. A soft-max layer produces class scores ("Action Scores") representing driver-state categories such as "Drive both hands," "Drink," "Call," and "Text." EX1006, ¶¶[0101], [0113]. Chao's architecture, comprising a CNN backbone feeding a classification head, corresponds directly to a distraction-classification prediction head consuming backbone features and outputting categorical labels. EX1003, ¶¶120-125.

c. Motivation to Combine (MTC)

A POSITA would have viewed Chun-Chao as textbook candidates for combination under *KSR*: Chun already provides the exact modular backbone-plus-heads framework into which Chao's distraction-classification head naturally plugs. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 417 (2007). Chun proves, in practice, that its FPN backbone can simultaneously feed multiple task-specific heads (pose, seat-belt). Chao teaches a distraction-classification head that consumes the same type of backbone-generated feature maps and outputs categorical distraction state labels. Each reference performs its known function. The combination is simply placing a known distraction-classification head onto a known multi-head backbone that was architected for exactly that type of modular expansion.

The how is obvious. Both references are built on standard multi-task CNN conventions (e.g., shared backbone → specialized heads). Chun already

demonstrates the architectural pattern of branching multiple heads off the same feature maps. Chao's distraction-classification module is implemented in the same way, taking feature maps as inputs and emitting a label. Attaching Chao's classification head to Chun's existing FPN backbone would require no novel architecture, no change to feature extraction, and no redesign of the loss pipeline. A POSITA would connect Chao's head at the same point where Chun attaches its pose/seat-belt heads.

The why is also obvious. Chun expressly identifies distracted-driving behaviors (drinking, phone-use, yawning, leaning on the door) as the problem being addressed. A POSITA would immediately understand that classifying such behaviors into distraction categories is merely the natural next step, i.e., converting the pose/features Chun already extracts into an actual distraction classification. That is exactly what Chao provides. Thus, the motivation is embedded in Chun's own problem statement, not hindsight.

The reasonable expectation of success follows directly. By the priority date, adding a classification head onto a shared FPN backbone was a routine multi-task extension (e.g., Mask-R-CNN and its variants). Chun already proves the backbone can drive multiple heads concurrently. Chao demonstrates a distraction-classification head that uses the same type of backbone features (face/hand/body-related) that Chun already extracts. A POSITA would fully expect that (i) training

the branch, and (ii) receiving categorical distraction labels would work the first time, because both sides of the interface (shared features → classification head) were already conventional and known to interoperate.

There were no technical obstacles. Both use the same conventional CNN primitives. Both operate in the same in-vehicle human-state-inference field. Both aim to solve the same overarching objective, i.e., detect driver distraction. Under *KSR*, that is the exact scenario where the Supreme Court says the combination is obvious: a predictable use of known components, each performing its established function, to yield a straight-line improvement that a POSITA would have been motivated to make and would have reasonably expected to succeed in making.

Chun-Chao renders obvious the distraction-classification prediction head . EX1003, ¶¶126-136.

5. [1E]

Chun discloses [1E]. Chun teaches an object-detection head for seat-belt identification. Chun explains IVMS detect "correct donning of a seat belt" using vision-based sensing and cites SSD and other detectors. EX1005, 1-3. NADS-Net includes a seat-belt prediction head that "produces a probability density function… indicating the likelihood of a pixel being a seat belt," which is thresholded to a binary mask. EX1005, 7.

Figure 5 shows an example of this seatbelt detection:



Figure 5: Seat belt detection result (*left*) in comparison with the ground truth (*right*). With human inspection, the prediction result is of good quality as it correctly marks the seat belt area. However, IOU for this particular example was 46%, justifying the need for a better evaluation metric.

EX1005, FIG. 5.

A POSITA would recognize the seat-belt head performs object detection: it identifies the "seat belt" class and its location via a dense probability map, functionally analogous to bounding-box detection in Faster R-CNN or SSD cited by Chun. EX1003, ¶¶137-141.

Chun discloses an object-detection prediction head satisfying [1E].

6. [1F]

Chun discloses [1F]. Chun teaches pose estimation heads within NADS Net. Chun states the key point and part affinity field heads together perform "2D driver/passenger pose estimation." EX1005, 1. Chun collected in vehicle images of drivers performing behaviors such as "drinking, using smartphones, yawning, sneezing, leaning on the door, adjusting the sun visor" to supplement MS COCO. EX1005, 4. The key point head predicts pixel probabilities for each joint and the

PAF head encodes pairwise joint relationships and the "outputs from the first two heads are used to parse the key point instances into human skeletons." EX1005, 5.

Chun discloses pose estimation prediction heads of [1F]. EX1003, ¶¶142-145.

7. [1G]

For avoidance of doubt, Petitioner does not contend [1G] requires a "joint loss"; the plain language is satisfied so long as a loss "based on outputs of the one or more prediction heads" is minimized in a way that updates the backbone parameters, including where that loss is derived from a single prediction head. Nonetheless, regardless of interpretation, Chun-Chao renders obvious [1G]. EX1003, ¶146-147.

a. Chun

Although Chun does not expressly state "loss minimization," its TensorFlow multi-head CNN, supervised training, and accuracy/transfer-learning experiments indicate standard loss minimization and backpropagation. EX1003, ¶¶148-149. Chun trains NADS-Net on labeled driver images using key-point, PAF, and seat-belt heads sharing an FPN backbone, reports "training and validation" across datasets, and compares "pretrained" versus "transferred" models, EX1005, 5-7, terminology a POSITA would understand as iterative joint-loss minimization updating shared backbone parameters. EX1003, ¶149. Each head outputs a

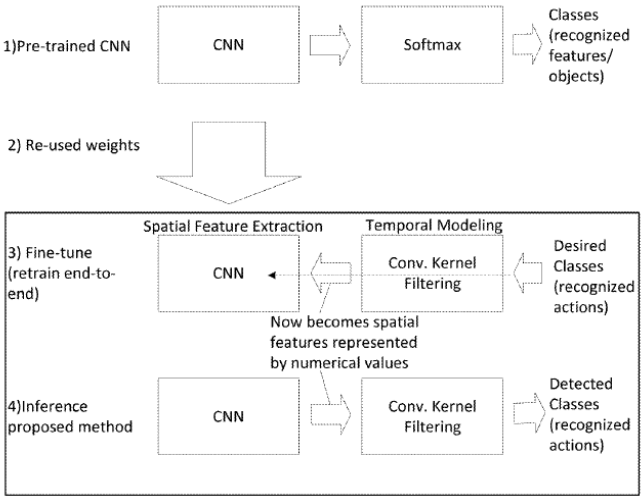
probability map that is compared to ground truth, and the combined task losses are minimized jointly; the conventional multi-task approach used in OpenPose and Mask R-CNN, which Chun cites.

A POSITA would understand Chun to be minimizing a joint loss across the prediction-head outputs to update the shared backbone, and would look to Chao for the implementation details of that joint-loss training. EX1003, ¶¶150-154.

b. Chao

Chao expressly teaches end-to-end supervised training. Chao states the system uses "objective functions (i.e., loss functions)" and "any kind of training principle, such as backpropagation." EX1006, ¶¶[0041], [0046]. Chao explains that prediction outputs are used "to further train the CNN." EX1006, ¶[0098]. Figure 8A confirms that the "pre-trained CNN" is then "fine-tune[d] (retrain end-to-end)," meaning the backbone parameters are updated based on final output loss.

FIG 8A



EX1006, FIG. 8A.

Chao teaches backbone parameters are updated based on final classification outputs. Chao states the "2D CNN 802 may change the output … to represent a plurality of values for the new desired 'action classes.'" EX1006, ¶[0099]. Chao explains "feedback from the results produced by [the] ST system" is used "to train the CNN" with "backpropagation." EX1006, ¶¶[0046], [0098]. A POSITA would view this as standard supervised learning in which prediction-head losses backpropagate to update the backbone. EX1003, ¶¶155-162.

c. MTC

A POSITA would recognize that applying Chao's loss-minimization to Chun's multi-head CNN is simply standard practice. The norm was to jointly optimize backbone and heads by minimizing task-output loss, and Chao says so

explicitly. Chun suggests this conventional multitask training, even if it omits the loss formula. A POSITA would therefore fill that gap with Chao's explicit teaching because it is the ordinary way to improve CNN accuracy in this field.

A POSITA reading Chun-Chao together would immediately see how to combine them and why it improves accuracy. Chun already uses a shared backbone with multiple heads for driver monitoring and trains them end-to-end in TensorFlow. Chao supplies the missing explicit piece: a loss-minimization routine in which the distraction-classification loss is backpropagated to update those shared backbone weights.

The how is the same standard workflow used before the '243 patent: define a classification loss for the distraction head and let the framework backpropagate it into the shared FPN weights. No architectural change is needed.

The why is straightforward. Chun seeks optimal driver-state accuracy, and loss minimization is the known way to improve it. Chao tells exactly how to implement that loss. A POSITA would apply that routine to Chun because it is the standard and obvious method to optimize weights for distraction recognition and Chun is already structured to benefit

A POSITA would also have held a strong expectation of success. At the time, loss minimization with gradient based weight updates was the default paradigm in deep learning, and TensorFlow already propagated gradients from the

head back into the backbone automatically. No re-architecting or new training logic is required, because the approach is simply the ordinary application of an already established technique to the same architecture and the same task. A POSITA would reasonably expect that adding Chao's explicit loss to Chun's backbone and heads would predictably improve distraction-classification accuracy in the same driver monitoring context.

[1G] represents nothing more than a predictable substitution. EX1003, ¶¶163-171.

8. [1H]

Chun-Chao renders obvious [1H].

a. Chun

Chun teaches a trained, deployable CNN ("more nimble and thus more suitable for in-vehicle monitoring" and "non-invasive, non-obstructive means to monitor and detect in-cabin activities") which a POSITA would understand requires storing the trained network weights (backbone + heads) for runtime inference. Chun's discussion of comparing "trained only on MS COCO," "trained only on the new driving dataset," and "pretrained then transferred," EX1005, 7, strongly suggests checkpointing finalized models for evaluation and deployment. EX1003, ¶¶173-175. Where Chun suggests it, Chao is explicit.

b. Chao

Chao states "a memory of the one or more memories 104 may store … a trained system, such as a neural network, or a deep neural network." EX1006, ¶[0054]. Chao's disclosed end-to-end training loop, EX1006, ¶¶[0041], [0046], [0098], confirms that after gradient-based updating, the trained network is saved in memory for inference. EX1003, ¶176

c. MTC

A POSITA reading Chun-Chao together would understand how and why to store the trained distraction model. Chun already trains multiple heads and a shared backbone until convergence for in vehicle inference. Chao provides the explicit teaching that once training is complete, learned backbone and head parameters are saved to non-volatile memory for later loading. After convergence, the framework storage primitive is invoked to serialize those parameters to a file or checkpoint, which TensorFlow and PyTorch already provide. Chun must store trained parameters to run in a vehicle, so adding Chao's storage step is the ordinary final deployment step.

A POSITA would reasonably expect storing the fully trained model to work on first application because serialization and re-loading were routine. Saving optimized backbone weights and prediction-head weights as a single artifact was the standard way to deploy a trained CNN. Nothing in Chun's architecture or data

flow prevents storage, and nothing in Chao is incompatible with Chun. Both references assume the same CNN paradigm, the same backbone plus heads, and the same field of use. The combination simply applies the known deployment step to the known multi-head model, yielding the predictable result: the same model Chun already trains is now persisted and ready for inference in the vehicle runtime environment. Under *KSR*, that is obvious. It is the ordinary use of a known technique model storage to improve a closely related system that is already designed for deployment. *KSR*, 550 U.S. at 417; EX1003, ¶¶177-183.

B. Claim 2

Chun-Chao renders claim 2 obvious. EX1003, ¶184.

1. Chun

Chun expressly discloses prediction heads made of multiple convolutional layers coupled to a backbone (two 3×3 convolutions and a 1×1 convolution) that output pixel-wise probability maps for key-point, part-affinity, and seat-belt heads. EX1005, 5, FIG. 3; EX1003, ¶¶185-186. Claim 1 shows Chun-Chao renders obvious the distraction-classification head. EX1003, ¶186. Although Chun does not list pooling or fully connected layers, a POSITA would understand convolution plus pooling plus fully connected layers was the standard classifier-head pattern. Adding that branch for distraction classification to Chun's FPN features would have been obvious, as taught in Chao. EX1003, ¶187.

2. Chao

Chao's ST head receives features "extracted from the 2D CNN model" and applies convolutional layers for feature refinement. EX1006, ¶¶[0099], [0106]; EX1003, ¶¶188-190. Chao discloses dimensionality reduction before the fully connected layer (Nk < Ns), which a POSITA would recognize as the pooling step. EX1006, ¶[0109]. Chao then applies a fully connected layer that outputs "action scores" identifying distraction states. EX1006, ¶[0107]; EX1003, ¶¶191-192.

3. MTC

A POSITA would understand Chun already supports adding new task heads, and Chao supplies the standard conv then pool then FC pattern for classification heads. The classifier would attach to Chun's shared FPN features so convolution refines features, pooling aggregates them, and the FC layer outputs distraction class scores. No redesign is required. Both references target driver state recognition and use the same modular backbone-plus-heads paradigm. Conv then pool then FC was already the conventional classification head pattern, and adding it to Chun would yield categorical distraction labels using the same backbone features.

A POSITA would have had a strong reasonable expectation of success. Chun's backbone encodes the exact cues Chao uses: face geometry, hand motion, body pose. Attaching a conv-pool-FC branch to those features was a routine extension that would predictably yield accurate categorical distraction predictions.

This is the familiar application of known components in the familiar way to yield a familiar result, exactly the type of combination *KSR* identifies as obvious. *KSR*, 550 U.S. at 417, 421; EX1003, ¶¶193-197.

C. Claim 6

Petitioner is not construing Claim 6 to require that the distraction-classification model must include the "distraction classification prediction head" or that the distraction-classification prediction head produces the distraction-classification tag. Claim 6 only recites that the model receive image data and output a distraction-classification tag, and that is satisfied whether the model uses a dedicated classification head, a shared head, or a single prediction head that performs the classification function. Regardless of the interpretation, Chun-Chao renders obvious claim 6. EX1003, ¶198-199

1. Chun

Chun teaches inputting in-vehicle driver images into a multi-task CNN and producing outputs such as pose estimation and seat-belt detection. EX1005, 1, 4-5. Chun shows training on annotated distraction behaviors such as drinking and phone use, and a POSITA would understand the same trained network receives real-time driver images and outputs predictions during inference. EX1005, 4-5.

Adding Chao's classification head to Chun would therefore yield the recited step of inputting image data into a distraction-classification model and receiving a distraction-classification tag. EX1003, ¶¶200-201.

2. Chao

Chao expressly discloses inputting driver images into a CNN based driver monitoring system and receiving a categorical output identifying driver distraction state. EX1006, ¶¶0129, 0136. Chao's architecture uses a shared 2D CNN backbone and a SoftMax classification layer that emits "action scores." EX1006, ¶0107. Chao identifies concrete distraction classes such as "Drink," "Call," "Text," "Makeup," and binary tags such as "FULL_ATTENTION" versus "DRIVER_TEXTING." EX1006, ¶[0113], FIG. 6. A POSITA would understand Chao discloses receiving a discrete distraction-classification tag based on image data. EX1003, ¶¶202-206.

3. MTC

A POSITA would have added Chao's classifier branch to Chun's existing multi-head CNN. Chun already routes driver-image frames through a shared FPN backbone into multiple heads for pose and seat-belt detection, and Chao provides a standard SoftMax classifier head that consumes the same type of backbone features and outputs a categorical distraction label. Attaching Chao's head to Chun's shared feature maps would yield the claimed distraction-classification tag. That is the

same conv-backbone-plus-classification-head pattern that was industry standard before the '243 patent.

A POSITA would also know why. Chun seeks improved distracted driving detection, and Chao supplies the class labels that make the detection actionable. Combining them yields the expected real time interpretable output without altering the backbone or training pipeline. Success would have been expected because both references already run real time inference using the same CNN conventions. Attaching a standard SoftMax branch to existing FPN features was a routine multitask extension requiring no new architecture. Under *KSR* this is the predictable combination of familiar s yielding a predictable result. EX1003, ¶¶207-218.

D. Claims 7 and 8

Chun-Chao renders obvious claims 7 and 8. EX1003, ¶¶219-220, 232-233

1. Chun

A POSITA would understand Chun-Chao discloses generating a distraction-classification tag identifying states such as texting or calling. EX1003, ¶¶221, 234. It would have been obvious to then perform a downstream action based on that tag. Chun states IVMS "play a critical role in preventing and mitigating traffic accidents by alerting the distracted driver and adaptively adjusting the safety mechanisms." EX1005, 1; EX1003 ¶¶222, 235. A POSITA would view Chun's

CNN outputs as the perception signals that drive those alerts and adjustments, as taught in Chao.

2. Chao

Chao makes the actuation explicit: after determining the driver's state, the "spatio-temporal analysis modeling scheme" "trigger[s] vehicular actions based on the determined state." EX1006, ¶¶[0022], [0129]; EX1003 ¶¶223, 236. Those actions include "modifying … acceleration, braking, [or] steering" or "communicating … notifications via the UIs." EX1006, ¶[0095]. Notifications may be "visual," "sound," or "vibration." EX1006, ¶[0168]; EX1003, ¶¶224, 237. A POSITA would equate "visual" with updating a display and "sound" with an audible alarm; and would view forwarding notifications to other processors, EX1006, ¶¶[0117], [0032], as transmitting the prediction to a remote controller/server. EX1003, ¶¶224, 238-240.

3. MTC

A POSITA would have combined Chun-Chao by routing the distraction-classification tag output from Chao's classifier head (integrated into Chun's multi-head CNN) to the control module disclosed in Chao. EX1006, ¶¶[0095]-[0168]; EX1003, ¶¶225, 241. At runtime, Chun's backbone receives driver images, extracts shared features, and those features are fed to Chao's classifier, which outputs a categorical distraction tag. EX1003, ¶¶225, 241. That tag then conditions the driver

alert or control adjustment using routine IVMS message-passing between the CNN and the controller. EX1003, ¶¶225, 241-242.

Chun already extracts features for distracted behaviors such as "drinking, using smart phones, exercising, yawning, sneezing, leaning on the door, putting hands out the window, [and] adjusting sun visor." EX1005, 4; EX1003, ¶¶226, 242. Chao identifies nearly the same classes (e.g., "Drink," "Call," "Look right," "Read," "Text"). EX1006, ¶[0113]; EX1003 ¶¶226, 242. A POSITA would view adding Chao's ST classification branch to Chun's backbone as a direct modular extension because both use shared backbone features feeding prediction heads. Chun states IVMS "alert the distracted driver and adaptively adjust the safety mechanisms." EX1005, 1. Chao discloses triggering vehicular actions based on state. EX1006, ¶[0065]. Combining them yields a predictable result, actionable distraction tags, with expected success. EX1003, ¶¶227-231, 243-245.

E. Claim 9

Chun-Chao renders obvious claim 9. EX1003, ¶246

Chun already provides the shared-backbone architecture generating feature maps for multiple heads. EX1003, ¶¶247-250.

Chao: The claim simply recites that the distraction-classification head outputs a label (undistracted / mobile-usage / other distracted), as Chao already discloses. Chao teaches the ST output is "a vector of NC values" normalized by

SoftMax to form a "probability distribution vector," and that "each of the classes of the vector" is output as the driver-state label ("Action Score"). EX1006, ¶[0101]; EX1003 ¶¶251-252. Table 1 maps those class labels into categories a POSITA would recognize as undistracted (e.g., "Drive both hands"), mobile usage (e.g., "Call," "Text"), or other distracted (e.g., "Adjust radio," "Makeup," "Reach side"). EX1006, ¶[0113]; EX1003 ¶¶253-255.

TABLE 1

Exemplary list of classes (i.e. states) and corresponding reaction time values

	Index												
	1	2	3	4	5	6	7	8	9	10	11	12	13
Class	Drive both hands	Switch gear	Talk to passenger	Adjust Radio	Drink	Call left hand	Call right hand	Hair makeup	Reach side	Reach behind	Text left hand	Text right hand	Hands free
Reaction time	2.0 sec	2.1 sec	2.4 sec	3 sec	3 sec	3.2 sec	3.2 sec	3.5 sec	4 sec	4 sec	4.2 sec	4.2 sec	5 sec

EX1006, Table 1.

Chao supplies the explicit SoftMax classification head that consumes those maps and outputs categorical labels. EX1003, ¶256.

MTC: For the reasons given for claims 6-7, a POSITA would have found it obvious to integrate Chao's classifier branch into Chun's backbone to output the distraction-state labels recited in claim 9. EX1003 ¶¶256-261.

F. Claim 10

Chun discloses claim 10. EX1003, ¶262. NADS-Net uses a shared FPN backbone with multiple heads, including a seat-belt head that outputs a pixel-probability map thresholded to a binary mask. EX1005, 1-2,7; EX1003, ¶¶263-264. Chun states this provides "more granularity in detection results" to locate the belt. EX1005, 2; EX1003, ¶264. Chun teaches all heads use the same backbone feature vector. EX1005, 2, 5; EX1003 ¶264. A POSITA would understand the seat-belt head receives that backbone vector as its input and performs object-location detection.

Chun discloses the claimed object-detection prediction head. EX1003, ¶264.

G. Claim 11

Chun discloses claim 11. EX1003, ¶265. Chun's FPN backbone feeds prediction heads for 2D pose estimation, including key-point and PAF heads, trained on distracted-driving gestures such as drinking, phone use, yawning, leaning, and adjusting the sun visor. EX1005, 1, 4; EX1003, ¶266. Chun states the PAF head outputs vector fields capturing pairwise joint relationships. EX1005, 2, 5; EX1003 ¶267. A POSITA would understand the PAF head receives the same backbone feature vector and uses that vector to infer joint orientation. EX1003, ¶267.

Chun discloses a pose-estimation head that detects orientation of body parts from an input vector. EX1003, ¶268.

H. Claim 12

Chun-Chao renders this claim obvious. EX1003, ¶269

Chun presents NADS-Net as an IVMS-deployment model "more nimble and thus more suitable for in-vehicle monitoring" and capable of "non-invasive, non-obstructive" cabin monitoring. EX1005, 1; EX1003, ¶270. Chun's deployment teaching implies storing trained weights for inference in-vehicle, and shows cabin-mounted cameras, including one "mounted below the rear view mirror" with dashboard IR illumination. EX1005, 4; EX1003 ¶271. Chun suggests deployed storage proximate to dash-region sensing hardware, but may be alleged to be silent on storing in the dash-mounted camera itself, which is explicit in Chao. EX1003 ¶272.

Chao makes that explicit: memory 104 "may store … a trained system, such as a neural network," and the camera acquisition device may be "on the dash." EX1006, ¶¶[0054], [0086], FIG. 5A; EX1003, ¶¶273-275. Chao teaches co-locating camera, processor, and memory in one housing. EX1006, ¶[0258]; EX1003 ¶276.

MTC: A POSITA would have viewed storing the trained model in the dash mounted camera as a routine deployment choice suggested by Chao and consistent

with Chun's deployment context. EX1003 ¶¶277-278. A POSITA would load the trained NADS Net weights into the local memory of the dash mounted camera hardware that Chao discloses. Chun-Chao use CNNs for in vehicle driver monitoring with interior cameras, and storing trained weights in the same housing as the capture device was already routine. Applying Chao's deployment architecture to Chun's trained model is a straightforward engineering step. Chun's goal is real-time distracted driving detection, and Chao discloses storing the trained network in the dash-mounted camera for in vehicle inference. A POSITA would expect success because both references use the same backbone-plus-head CNN paradigm, both already perform real-time inference, and storage of trained weights is task agnostic. EX1003, ¶¶279-281.

I. Claims 13, 14, 18-20, 24, and 25

Claims 13, 14, 18-20, 24, and 25 add no new features. They restate the limitations of claims 1, 2, 6-9, 11, and 12 as computer-readable medium (CRM) or device. EX1003, ¶282. Claim 13 mirrors claim 1 as a non-transitory medium; claims 14, 18, and 19 mirror claims 2, 6, and 7; claim 20 mirrors claim 1 as a processor-and-memory device; and claims 24 and 25 mirror the same refinements of claims 6, 7, 18, and 19. EX1003, ¶¶283-284. A POSITA would understand that the same Chun-Chao method is stored as program instructions or executed on in-vehicle processors. Chao expressly discloses memory storing a trained neural

network and processors running it. EX1006, ¶¶[0092]-[0093]; EX1003, ¶286. Converting an obvious method into standard software or device form would have been routine. For the same reasons as claims 1, 2, and 6-12, these claims are likewise obvious. EX1003 ¶¶285-288.

J. Independent Claim 26

Claim 26 recites the same subject matter as claim 1, with the same additional features already addressed for claims 6, 7, 8, and 12. EX1003, ¶289. Each component limitation has already been shown obvious in Chun-Chao. EX1003, ¶289. Claim 26[A] is the same dash-mounted storage limitation as claim 12: Chun suggests in-vehicle model deployment and Chao expressly discloses storing the trained network in the dash-mounted camera. EX1003, ¶290. Claim 26[B] is the same backbone-plus-heads architecture and supervised training as claim 1, already taught by Chun-Chao. EX1003, ¶291. Claim 26[C] recites capturing an interior vehicle image using a dash-mounted sensor, which both Chun-Chao teach. EX1003 ¶¶292-294. Claim 26[D] is the same image input and distraction tag output as claim 6, and claim 26[E] is the same downstream action as claims 7 and 8. EX1003 ¶¶295-296. For the same reasons as claims 1, 6-8, and 12, claim 26 is obvious in view of Chun-Chao. EX1003, ¶289.

K. Claim 27

Claim 27 recites substantially the same subject matter as claim 9. Chun-Chao discloses classifying the driver image into categorical labels corresponding to undistracted, mobile usage, or other distracted states using an input vector generated based on images of drivers. EX1003, ¶297.

L. Claim 28

Chun discloses claim 28. Chun's seat belt detection head produces a probability density over the image domain "indicating the likelihood of a pixel being a seat belt," which is then thresholded to produce a binary segmentation mask that identifies the seat belt location relative to the driver. EX1005, 3, 7, FIG. 5; EX1003, ¶298. An example of this seatbelt detection is in Chun Figure 5:



Figure 5: Seat belt detection result (*left*) in comparison with the ground truth (*right*). With human inspection, the prediction result is of good quality as it correctly marks the seat belt area. However, IOU for this particular example was 46%, justifying the need for a better evaluation metric.

EX1005, FIG. 5.

This seat belt detection prediction head provides the exact location of a seatbelt in an image of a driver/passenger using the pixels of the image. A

POSITA would understand this as detecting the location of an object within the image. EX1003, ¶299.

M. Claim 29

Chun describes a feature pyramid backbone with prediction heads for "2D driver / passenger pose estimation," including key point and PAF heads trained on images of drivers performing gestures such as "drinking, using smart phones, yawning, sneezing, leaning on the door, putting hands out the window, searching floor and the center console, adjusting sun visor, and etc." EX1005, 1, 4; EX1003, ¶300.

Chun explains that the PAF head produces vector fields that "encode pairwise relationships between body joints" indicating how keypoints are connected. EX1005, 2, 5; EX1003 ¶301. A POSITA would understand that the PAF head detects body part orientation depicted in the image. EX1003 ¶301.

Chun discloses a pose estimation prediction head, detecting orientation of one or more body parts in image data. EX1003, ¶302.

N. Claims 30-33

Claims 30-33 restate the same features of claims 26-29 in device form rather than method. EX1003, ¶303. Claim 30 mirrors claim 26, but recites a dash-mounted camera device with processor and memory, and claims 31-33 mirror claims 27-29 (label classification, object-location detection, and pose-orientation

detection). EX1003, ¶304. A POSITA would understand that the same distraction-classification functionality taught by Chun-Chao is executed by processors in an in-vehicle device. EX1003, ¶305. Chao expressly discloses such modules, including in-vehicle cameras, memory storing a trained neural network, and processors. EX1005, ¶¶[0092]-[0093]; EX1003, ¶306. Recasting the same obvious operations into "device" form adds no inventive distinction. For the same reasons as claims 26-29, claims 30-33 are obvious in view of Chun-Chao. EX1003, ¶307.

VI. GROUND 2: CHUN IN VIEW OF CHAO AND FURTHER IN VIEW OF SHANMUGAMANI RENDERS OBVIOUS CLAIMS 3-5, 10, 15-17, AND 21-23.

Chun in view of Chao and Shanmugamani discloses each added feature of dependent claims 3-5, 10, 15-17, and 21-23. These claims recite routine CNN blocks already used in Ground 1 (convolution stacks, pooling, fully connected layers, bounding box regression, class heads, and pose heads), which were standard TensorFlow/Keras components for object detection and image classification. EX1003 ¶¶308-309.

A POSITA would have combined Chun, Chao, and Shanmugamani because Chun expressly states NADS-Net was implemented in TensorFlow, and Shanmugamani teaches the same off-the-shelf TensorFlow layer patterns for classification, detection, and pose estimation used by Chun and Chao. EX1003, ¶¶310-317. This is not inventive: Shanmugamani provides step-by-step

TensorFlow CNN implementations for autonomous-camera workloads and confirms that the dependent-claim details reflect conventional blocks. EX1003, ¶¶318-319.

A POSITA would have had a reasonable expectation of success, and dependent claims 3-5, 15-17, and 21-23 are obvious for the same reasons as Ground 1. EX1003, ¶¶320-323.

A. Claim 3

Chun-Chao-Shanmugamani renders obvious claim 3 that depends from claim 1. EX1003, ¶324. Chun-Chao renders claim 1 obvious. EX1003, ¶325. Claim 3 simply recites that the object detection head includes a bounding-box regression network to output box coordinates and a class-prediction network to output an object label. EX1003, ¶325.

1. Chun-Chao

Chun already teaches object localization. Chun states the seatbelt head provides "proper/improper use cases judged by the relative position of seat belt to the detected body position." EX1005, 3. This is object localization of the same type used in bounding box detection as in Faster R-CNN and SSD which Chun cites. EX1003 ¶¶326-327.

Chao likewise teaches determining an object footprint and classifying it. EX1006, ¶¶[0046], [0075], [0083].

A POSITA would therefore understand Chun-Chao's seatbelt head as implemented through bounding box regression followed by class prediction, and Shanmugamani provides explicit disclosure of the dual head object detection structure recited in claim 3. EX1003, ¶¶328-329.

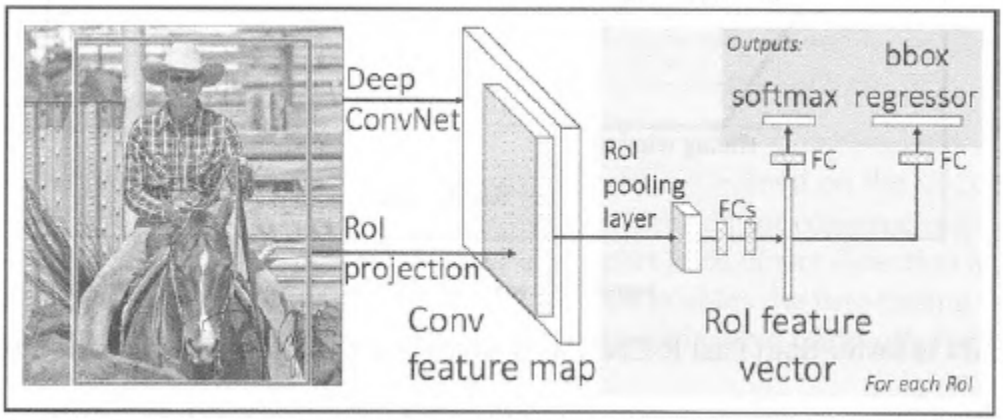
2. Shanmugamani

Shanmugamani explains CNNs are "a collection of perceptrons and activation functions" and "CNNs have weights, biases, and outputs through nonlinear activation" trained to "detect features or patterns in images." EX1010, 7, 18; EX1003, ¶330

For object detection, Shanmugamani states "object detection is the act of finding the location of an object in an image" by finding objects "along with the labels." EX1010, 109-110. Shanmugamani describes localization as a regression problem where "the bounding box is four numbers" and that there is also a "need to predict the label which is a classification problem." EX1010, 117. Shanmugamani confirms "the regression encoder and classification encoder function independently" and a bounding box may be parameterized by "center of the coordinates with the height and width of the bounding box." EX1010, 119; EX1003, ¶331.

Shanmugamani then provides concrete TensorFlow examples of this exact dual head structure, including R CNN, Fast R CNN, Faster R CNN, SSD, and

YOLO which all output bounding box coordinates and class predictions. EX1010, 120-123. For example, in R CNN, a "linear regression model to predict a closer bounding box is trained with object region proposals" and "the last layer of the CNN is trained with an SVM identifying the object with a no object class." EX1010, 120. Fast R CNN uses ROI pooling with "one hot encoding with softmax" and a bounding box regressor. EX1010, 121. SSD "simultaneously predicts the object and finds the bounding box." EX1010, 122-123; EX1003 ¶¶332-336.



EX1010, 121.

These teachings show a bounding box regression network plus a classification network is the standard implementation of object detection in contemporary CNNs. EX1003, ¶¶337-338. A POSITA therefore would understand that claim 3 recites the well-known dual head pattern taught directly in Shanmugamani, including for automotive computer vision. EX1003, ¶339.

3. MTC

A POSITA would have been motivated to combine Chun-Chao-Shanmugamani. EX1003, ¶340. Chun already implements NADS-Net in TensorFlow, uses COCO, and detects seat belts, but does not spell out the standard TensorFlow layer recipe for bounding-box, regression-plus-class prediction. A POSITA using Chun's TensorFlow platform would naturally consult a contemporaneous TensorFlow reference, and Shanmugamani provides exactly that: explicit TensorFlow object-detection implementations using COCO, with bounding-box regressors and class-prediction heads trained end-to-end. EX1003, ¶340.

Chun-Chao already identify seat-belt location and classify object state relative to the driver. Implementing that in the standard TensorFlow form taught by Shanmugamani would require no redesign; it is simply the routine COCO-style bounding-box plus class-prediction pattern. EX1003, ¶340.

And Shanmugamani expressly states the reader "will learn how to train deep learning models for computer vision applications and deploy them on multiple platforms," including autonomous-vehicle detection. *Id.*, 6. Because Chun already uses TensorFlow, a POSITA would expect straightforward success inserting Shanmugamani's standard TensorFlow object-detection blocks into Chun's TensorFlow seat-belt head. EX1003, ¶¶341-342.

Chun-Chao-Shanmugamani renders obvious claim 3. EX1003, ¶343.

B. Claim 4

Chun-Chao-Shanmugamani renders obvious claim 4. EX1003, ¶344. Chun-Chao-Shanmugamani renders obvious claim 3. EX1003, ¶345. Claim 4 only adds that the bounding box regression network and object class prediction network include hidden layers with convolution, batch normalization, and activation, which are conventional CNN s expressly taught in Shanmugamani.

1. Chun-Chao

Chun discloses multiple convolutional layers in its seat belt head ("two 3×3 convolutions and a 1×1 convolution"). EX1005, 5. A POSITA would understand these as hidden layers. Chao likewise teaches CNNs with input, output, and "one or more hidden layers." EX1006, ¶¶[0075], [0098]. While Chun-Chao do not enumerate the standard internal sub layers within each hidden layer, Shanmugamani provides those conventional implementation details. EX1003, ¶¶346-347.

2. Shanmugamani

Shanmugamani confirms that bounding box regression and object class prediction networks are conventionally implemented using convolutional, activation, and batch normalization layers. Shanmugamani teaches a neural network "is a collection of perceptrons and activation functions," where "inputs of x are passed through a hidden layer of perceptrons and summed to the output," and

"every layer can have one of the activation functions." EX1010, 11. Activation functions "make neural nets nonlinear" and adjust gradients during training, including sigmoid, tanh, and ReLU. EX1010, 7-9. Shanmugamani teaches batch normalization "increases the stability and performance of neural network training" because it "reduces overfitting and makes the network train faster" by "normalizing the output from a layer with zero mean and a standard deviation of 1." EX1010, 13.

Shanmugamani then walks through "retrain[ing] a model for our custom dataset," including dataset preparation, algorithm selection, and fine tuning. EX1010, 126. As an explicit example, Shanmugamani presents a YOLO based pedestrian detection model for self driving applications that includes convolutional layers, activation (tf.nn.leaky_relu), and regularization layers. EX1010, 128-130.

```
def pooling_layer(input_layer, pool_size=[2, 2], strides=2,
padding='valid'):
    layer = tf.layers.max_pooling2d(
        inputs=input_layer,
        pool_size=pool_size,
        strides=strides,
        padding=padding
    )
    add_variable_summary(layer, 'pooling')
    return layer
```

```
def convolution_layer(input_layer, filters, kernel_size=[3, 3],
padding='valid',
                      activation=tf.nn.leaky_relu):
    layer = tf.layers.conv2d(
        inputs=input_layer,
        filters=filters,
        kernel_size=kernel_size,
        activation=activation,
        padding=padding,
        weights_initializer=tf.truncated_normal_initializer(0.0, 0.01),
        weights_regularizer=tf.l2_regularizer(0.0005)
    )
    add_variable_summary(layer, 'convolution')
    return layer

def dense_layer(input_layer, units, activation=tf.nn.leaky_relu):
    layer = tf.layers.dense(
        inputs=input_layer,
        units=units,
        activation=activation,
        weights_initializer=tf.truncated_normal_initializer(0.0, 0.01),
        weights_regularizer=tf.l2_regularizer(0.0005)
    )
    add_variable_summary(layer, 'dense')
    return layer
```

```
yolo = tf.pad(images, np.array([[0, 0], [3, 3], [3, 3], [0, 0]]),
name='pad_1')
yolo = convolution_layer(yolo, 64, 7, 2)
yolo = pooling_layer(yolo, [2, 2], 2, 'same')
yolo = convolution_layer(yolo, 192, 3)
yolo = pooling_layer(yolo, 2, 'same')
yolo = convolution_layer(yolo, 128, 1)
yolo = convolution_layer(yolo, 256, 3)
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
yolo = pooling_layer(yolo, 2, 'same')
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
```

```
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
yolo = convolution_layer(yolo, 512, 1)
yolo = convolution_layer(yolo, 1024, 3)
yolo = pooling_layer(yolo, 2)
yolo = convolution_layer(yolo, 512, 1)
yolo = convolution_layer(yolo, 1024, 3)
yolo = convolution_layer(yolo, 512, 1)
yolo = convolution_layer(yolo, 1024, 3)
yolo = convolution_layer(yolo, 1024, 3)
yolo = tf.pad(yolo, np.array([[0, 0], [1, 1], [1, 1], [0, 0]]))
yolo = convolution_layer(yolo, 1024, 3, 2)
yolo = convolution_layer(yolo, 1024, 3)
yolo = convolution_layer(yolo, 1024, 3)
yolo = tf.transpose(yolo, [0, 3, 1, 2])
yolo = tf.layers.flatten(yolo)
yolo = dense_layer(yolo, 512)
yolo = dense_layer(yolo, 4096)

dropout_bool = tf.placeholder(tf.bool)
yolo = tf.layers.dropout(
        inputs=yolo,
        rate=0.4,
        training=dropout_bool
    )
yolo = dense_layer(yolo, output_size, None)
```

EX1010, 128-130.

Although the code snippet uses L2 regularization, Shanmugamani teaches overfitting is addressed by "regularizing using techniques such as dropout and batch normalization," which are standard TensorFlow functions. EX1010, 77. Thus, Shanmugamani expressly confirms that the conventional dual head object detection pattern of bounding box regression and class prediction is implemented with hidden layers that include convolution, activation, and batch normalization layers of the type claimed. EX1003, ¶¶348-351.

3. MTC

A POSITA would have combined Chun-Chao-Shanmugamani. Chun discloses a TensorFlow CNN for seat-belt detection, but does not enumerate the standard hidden-layer recipe. Chun states each head uses "two 3×3 convolutions and a 1×1 convolution" for pixel-wise outputs, EX1005, 5, but is silent on the activation / normalization sequencing. Chao likewise refers to neural networks with "one or more hidden layers" for vehicle-object classification. EX1006, ¶¶[0075], [0098]. Shanmugamani provides those missing implementation details by teaching standard TensorFlow object-detection heads built from convolution, activation, and batch-normalization layers. EX1003, ¶352.

A POSITA would have applied those conventional TensorFlow patterns to Chun's TensorFlow seat-belt head because Shanmugamani expressly teaches how to build and deploy such models for automotive perception and states the reader "will learn how to train deep learning models for computer vision applications and deploy them on multiple platforms." Shanmugamani p. 6. This is the predictable use of prior art according to their established functions, which is exactly the type of combination *KSR* identifies as obvious. *KSR*, 550 U.S. at 417; EX1003 ¶352.

A POSITA would have expected straightforward success. Chun-Chao-Shanmugamani renders obvious claim 4. EX1003, ¶¶353-356.

C. Claim 5

Chun-Chao-Shanmugamani renders obvious claim 5, depending from claim 1, and Chun-Chao renders claim 1 obvious. EX1003, ¶¶357-358.

1. Chun-Chao

Chun discloses pose estimation heads. Chun teaches keypoint and PAF heads that perform "2D driver/passenger pose estimation." EX1005, 1. Chun further states that each branch "employs two 3×3 convolutions and a 1×1 convolution to predict a pixel-wise probability distribution." EX1005, 5. These are TensorFlow implemented pose heads with multiple convolutional layers, but Chun does not spell out the conventional components within those hidden layers. Shanmugamani provides those missing TensorFlow details. EX1003, ¶359.

2. Shanmugamani

Shanmugamani teaches a neural network "is a collection of perceptrons and activation functions," where "inputs of x are passed through a hidden layer of perceptrons and summed to the output," and "every layer can have one of the activation functions." EX1010, 11. Activation functions "make neural nets nonlinear" and are used for gradient adjustment, including sigmoid, tanh, and ReLU. EX1010, 7-9. Shanmugamani teaches batch normalization which "increase[s] the stability and performance of neural network training" by "normaliz[ing] the output from a layer with zero mean and a standard deviation of 1." EX1010, 13.

Shanmugamani describes pose detection is "the act of finding joint locations in a human," by detecting "multiple locations such as head, neck, shoulders, ankles, and hands." EX1010, 117-118.



EX1010, 117.

Shanmugamani describes object localization in the same regression encoder and classification encoder pattern that independently outputs bounding box coordinates and labels. EX1010, 110, 114, 119. Shanmugamani teaches standard TensorFlow retraining on custom datasets including self-driving examples. The example code is as follows:

```
def pooling_layer(input_layer, pool_size=[2, 2], strides=2,
padding='valid'):
    layer = tf.layers.max_pooling2d(
        inputs=input_layer,
        pool_size=pool_size,
        strides=strides,
        padding=padding
    )
    add_variable_summary(layer, 'pooling')
    return layer
```

```
def convolution_layer(input_layer, filters, kernel_size=[3, 3],
padding='valid',
                      activation=tf.nn.leaky_relu):
    layer = tf.layers.conv2d(
        inputs=input_layer,
        filters=filters,
        kernel_size=kernel_size,
        activation=activation,
        padding=padding,
        weights_initializer=tf.truncated_normal_initializer(0.0, 0.01),
        weights_regularizer=tf.l2_regularizer(0.0005)
    )
    add_variable_summary(layer, 'convolution')
    return layer

def dense_layer(input_layer, units, activation=tf.nn.leaky_relu):
    layer = tf.layers.dense(
        inputs=input_layer,
        units=units,
        activation=activation,
        weights_initializer=tf.truncated_normal_initializer(0.0, 0.01),
        weights_regularizer=tf.l2_regularizer(0.0005)
    )
    add_variable_summary(layer, 'dense')
    return layer


yolo = tf.pad(images, np.array([[0, 0], [3, 3], [3, 3], [0, 0]]),
name='pad_1')
yolo = convolution_layer(yolo, 64, 7, 2)
yolo = pooling_layer(yolo, [2, 2], 2, 'same')
yolo = convolution_layer(yolo, 192, 3)
yolo = pooling_layer(yolo, 2, 'same')
yolo = convolution_layer(yolo, 128, 1)
yolo = convolution_layer(yolo, 256, 3)
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
yolo = pooling_layer(yolo, 2, 'same')
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
```

```
yolo = convolution_layer(yolo, 256, 1)
yolo = convolution_layer(yolo, 512, 3)
yolo = convolution_layer(yolo, 512, 1)
yolo = convolution_layer(yolo, 1024, 3)
yolo = pooling_layer(yolo, 2)
yolo = convolution_layer(yolo, 512, 1)
yolo = convolution_layer(yolo, 1024, 3)
yolo = convolution_layer(yolo, 512, 1)
yolo = convolution_layer(yolo, 1024, 3)
yolo = convolution_layer(yolo, 1024, 3)
yolo = tf.pad(yolo, np.array([[0, 0], [1, 1], [1, 1], [0, 0]]))
yolo = convolution_layer(yolo, 1024, 3, 2)
yolo = convolution_layer(yolo, 1024, 3)
yolo = convolution_layer(yolo, 1024, 3)
yolo = tf.transpose(yolo, [0, 3, 1, 2])
yolo = tf.layers.flatten(yolo)
yolo = dense_layer(yolo, 512)
yolo = dense_layer(yolo, 4096)

dropout_bool = tf.placeholder(tf.bool)
yolo = tf.layers.dropout(
        inputs=yolo,
        rate=0.4,
        training=dropout_bool
    )
yolo = dense_layer(yolo, output_size, None)
```

EX1010, 128-130.

In this example, the hidden layers include a convolutional layer, an activation layer (e.g., tf.nn.leaky_relu), and a regularization layer (e.g., L2 regularizer), and Shanmugamani explains overfitting may also be addressed using dropout or batch normalization. Shanmugamani further shows the output layer is fed by convolutional features and thus comprises a convolutional layer input. EX1010, 77; EX1003, ¶¶360-365.

3. MTC

A POSITA would have combined Chun-Chao-Shanmugamani for claim 5 for the same structural reason as claim 3, applied here to pose. Chun already discloses a TensorFlow-implemented pose-estimation head using "two 3×3 convolutions and a 1×1 convolution to predict a pixel-wise probability distribution." EX1005, 5. Chun therefore identifies the task (pose estimation) and framework (TensorFlow) but omits the standard sub-layer details in the hidden layers. Shanmugamani supplies those conventional TensorFlow layer components, teaching that a neural network "is a collection of perceptrons and activation functions" and that each hidden layer uses convolution, batch-normalization, and activation, and that the output head is also convolutional. EX1010, 6-13. Because Chun already uses TensorFlow for this prediction head, a POSITA would simply apply Shanmugamani's conventional TensorFlow pattern to implement Chun's pose head. This substitution would have been routine and predictable because Shanmugamani is expressly written as a TensorFlow implementation guide for CNN heads, including vision models for autonomous-vehicle workloads. EX1010, 6. Chun-Chao-Shanmugamani renders obvious claim 5. EX1003, ¶¶366-369.

D. Claims 15-17 and 21-23

Claims 15-17 and 21-23 recite the same features as claims 3-5 as CRM and device rather than method. Claim 15 is the CRM counterpart of claim 3, specifying

bounding-box regression and object-class prediction networks. Claim 16 corresponds to claim 4, reciting that those networks include hidden layers with convolution, batch normalization, and activation. Claim 17 corresponds to claim 5, reciting that the pose-estimation head includes hidden layers with convolution, batch normalization, activation, and a convolutional output. Likewise, claims 21-23 are the device counterparts of claims 3-5. Claim 21 corresponds to claim 3, claim 22 corresponds to claim 4, and claim 23 corresponds to claim 5, but expressed as logic executed by a processor with model parameters stored in memory. EX1003, ¶¶371-372.

The only differences are the statutory form (method vs. CRM vs. device). A POSITA would have understood that the same multi-task CNN disclosed and rendered obvious for claims 3-5 would be implemented either as software instructions on a non-transitory medium or as logic executed in a processor-based device. EX1003, ¶¶373-374.

Chao already discloses implementing distraction classification in software and hardware, including on-device processors, memory, and camera modules. EX1006, ¶¶[0092]-[0093]. Accordingly, for the same reasons set forth for claims 3-5, claims 15-17 and 21-23 would also have been obvious in view of Chun-Chao-Shanmugamani. EX1003, ¶¶375-376.

VII. GROUND 3: ZHENGYANG IN VIEW OF HE RENDERS OBVIOUS CLAIMS 1-6, 9-18, AND 20-24.

Zhengyang in view of He renders independent claims 1, 13 and 20 obvious. Zhengyang discloses the shared-backbone, multi-task CNN architecture for driver-state analysis. He provides the additional pose-estimation head that operates over the same shared feature maps. EX1003, ¶¶377-378. The dependent claims fall with the independents for the same reasons.

A POSITA would have been motivated to combine Zhengyang-He because both references address the same technical problem space of multitask visual human state recognition using a shared-backbone architecture and both provide complementary teachings that a skilled practitioner would have naturally integrated to improve driver state recognition performance. Zhengyang provides the foundation: a unified neural network model with a shared backbone and multiple task heads jointly trained through forward and backward propagation to "simultaneously recognize multiple features of the driver" and thereby "improve the accuracy of driver state recognition." EX1007, ¶¶[0003], [0005], [0046]. He then teaches precisely the type of additional branch a POSITA would have expected to attach to such a framework. He discloses a pose estimation "keypoint head" comprising convolutional and deconvolutional layers that predicts one hot keypoint masks for each joint type (for example left shoulder and right elbow) and that is trained jointly with detection and segmentation heads while sharing the

same backbone. EX1008, 7-8. The model in He "simultaneously predict[s] boxes, segments, and keypoints" in real time and consumes the same backbone feature maps for all heads. *Id.* A POSITA would have understood that pose estimation directly complements Zhengyang's distraction-classification tasks because keypoint geometry, such as head tilt, shoulder rotation, and hand position, are highly discriminative features for detecting whether the driver is distracted. Moreover, by the '243 patent priority date, adding a pose estimation head to a shared backbone multitask model was a routine and well-understood extension that required no change to the underlying training paradigm but merely the definition of another task head with its own labels and loss. Both references employ standard deep learning building blocks including ResNet style backbones, convolutional task heads, cross entropy losses, and gradient based optimization. A POSITA would have had a reasonable expectation of success in integrating He's keypoint head into Zhengyang's multitask system and would have expected predictable improvement in driver monitoring performance using ordinary skill rather than inventive ingenuity. EX1003, ¶¶379-385.

A. Independent Claim 1

1. [1Pre]

Zhengyang discloses a method by teaching acquiring driver images, inputting them into the multi-task network, generating feature maps, and outputting distraction-state results. EX1007, ¶¶[0020]-[0023], [0028], [0031], [0046]; EX1003, ¶¶386-387.

2. [1A]

Zhengyang discloses [1A]. Zhengyang describes supervised multi-task learning where annotated driver images are input into a neural network to train parameters. EX1003, ¶¶388-389. Zhengyang explains preprocessed driver images are input to the network "for attribute recognition" to determine whether the driver "is wearing a mask," "is wearing eyeglasses," "is smoking," or "is making a phone call." EX1007, ¶[0012]. Zhengyang makes clear this same structure is used during training: "before training… it is necessary to label the driver" and "the labeling information and driver images will be used as input to train the model." EX1007, ¶[0063]; *see also* EX1007, ¶[0025] ("labels the sample image data with multi-task attributes").

Zhengyang teaches inputting labeled examples into a neural network. EX1003, ¶¶390-393.

3. [1B] and [1C]

Zhengyang discloses both a backbone that extracts shared features from

driver images and task-specific prediction heads fed by those backbone features. Zhengyang describes a conventional supervised multi-task CNN pipeline in which "the collected RGB image is input into a backbone network to extract encoded features," EX1007, ¶[0060], and those shared feature maps are then routed to separate attribute-recognition modules. EX1003, ¶¶394-395.

Zhengyang identifies the backbone as SE-ResNet18: "the backbone network uses the SE-ResNet18 network." EX1007, ¶[0069]. The disclosure confirms how SE-ResNet18 is actually implemented: 112×112 crop, first 7×7→3×3, removal of initial pooling, stride modifications, BN / PReLU / dropout, EX1007, ¶[0100], e.g., standard refinements for cropped in-cabin faces. The 7×7 feature maps are then processed by OSME to extract features for "mask, eyeglasses, smoking, [and] phone-call" recognition. EX1007, ¶¶[0083], [0069]; EX1003, ¶¶396-400.

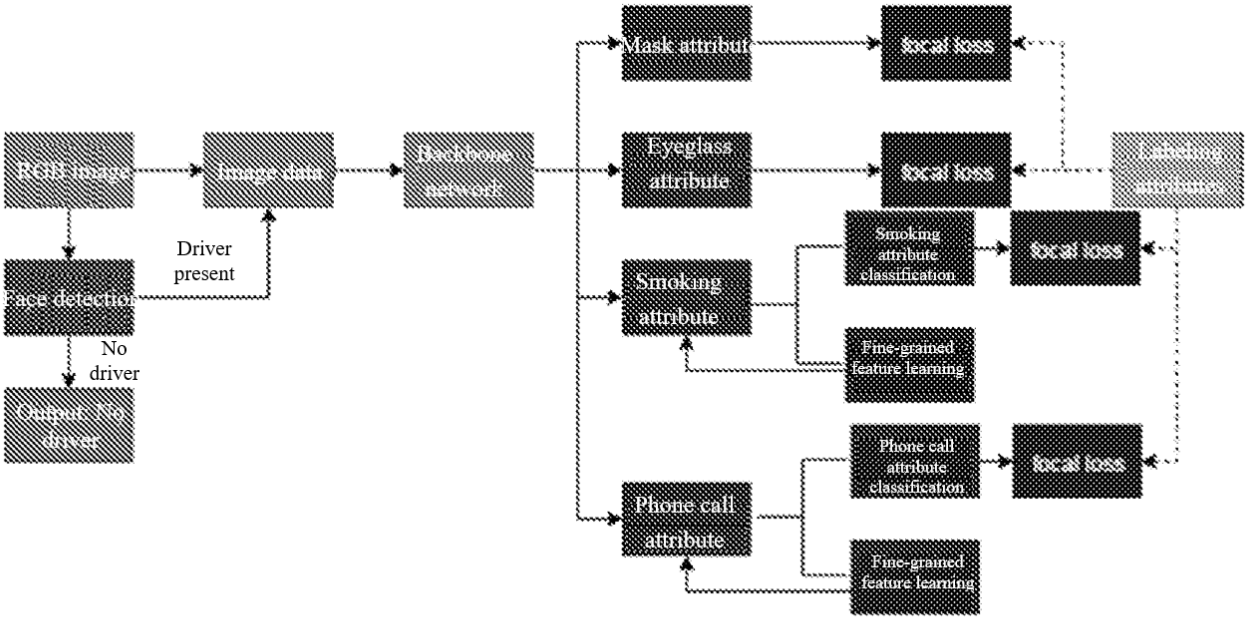


FIG. 3

EX1007, FIG. 3.

Zhengyang states these OSME-derived maps are "input into respective attribute recognition modules" which output "a feature map of each task and an attribute recognition result of each task." EX1007, ¶¶[0020]-[0021], [0109], FIG. 3. SE-ResNet18 the shared backbone and the attribute-recognition modules are one or more prediction heads coupled to that backbone. EX1003, ¶¶401-404.

4. [1D]

Zhengyang discloses [1D]. Zhengyang discloses prediction heads classifying smoking and phone use. Zhengyang states the multi-task network performs "mask recognition, eyeglasses recognition, smoking recognition, and phone call recognition," each as a separate downstream module receiving shared backbone

features. EX1007, ¶[0008]; *see also* EX1007, ¶¶[0020], [0094]. The outputs are binary, as Zhengyang specifies "the number of output categories … is two categories." EX1007, ¶[0028]; *see also* EX1007, ¶¶[0095], [0098]. The smoking- and phone-call-recognition modules are binary classification heads that determine whether the driver is engaged in a distraction behavior, i.e., analogous to the claimed distraction-classification prediction head. EX1003, ¶¶406-408.

Zhengyang discloses multiple backbone-coupled prediction heads, including smoking-and phone-call-recognition heads that perform distraction-classification. EX1003, ¶408.

5. [1E]

Zhengyang discloses [1E]. Zhengyang discloses an attention-proposal sub-network (APN) that performs object localization, i.e., the claimed object-detection prediction head. Zhengyang states the APN outputs "center point coordinates of the attention area and the edge lengths of the attention area" used for cropping the region of interest. EX1007, ¶¶[0033]-[0034]; *see also* EX1007, ¶[0101]. A POSITA would recognize this as bounding-box prediction serving the same role as region-proposal heads in two-stage detectors. EX1003, ¶¶409-410.

Zhengyang teaches iterative refinement: the cropped region "will be input into the first-level attention model … APN … will further crop … and input … into the second-level attention model … ." EX1007, ¶[0102]; *see also* EX1007,

¶[0103]. A POSITA would view this as hierarchical region-refinement consistent with standard detection practice. EX1003, ¶¶411-412.

Although described in conjunction with the attention module, the APN performs a distinct prediction function: producing coordinate-based bounding-box outputs derived from shared backbone features. EX1003, ¶¶412-413. A POSITA would understand the APN to be an object-detection prediction head separate from the distraction-classification heads within the disclosed multi-task architecture. EX1003, ¶413.

6. [1F]

Zhengyang-He renders obvious [1F].

a. Zhengyang

Zhengyang discloses classification and detection heads, but not a pose-estimation head. Adding such a head would have been obvious in view of He, given Zhengyang's shared SE-ResNet18 backbone producing dense spatial maps and multiple task-specific heads consuming those maps, plus attention-guided cropping capturing spatial cues suitable for pose analysis. EX1003, ¶¶414-416.

b. He

He teaches a pose-estimation prediction head, i.e., a keypoint branch predicting joint locations from shared features. He explains "Mask R-CNN… can easily be extended to human pose estimation," modeling each keypoint as a one-hot mask and predicting K masks (e.g., left shoulder, right elbow). EX1008, 7. A

POSITA would view this as a dedicated sub-network outputting spatial heatmaps. EX1003, ¶417. The keypoint head "consists of eight 3×3 512-d conv layers, followed by a deconv layer and 2× bilinear upscaling," trained using cross-entropy over an m^2-way softmax output, with accuracy measured using AP_{kp}. EX1008, 7-8. He confirms this branch operates in parallel with other heads: "We have a unified model that can simultaneously predict boxes, segments, and keypoints … ." EX1008, 8.

He expressly teaches a pose-estimation prediction head: an independent branch receiving shared convolutional features and producing spatial heatmaps for joints using a dedicated loss, which is readily integrable into Zhengyang's backbone-plus-heads design. EX1003, ¶¶418-419.

c. MTC

A POSITA would have been motivated to add He's pose-estimation head to Zhengyang's multi-task driver-monitoring backbone to improve driver-state assessment. Zhengyang already detects localized cues (smoking / phone use) but seeks to "comprehensively reflect the driver's driving state and improve the accuracy of driver state recognition." EX1007, ¶[0046]. He's keypoint head adds global pose (head/arm/body orientation) and is expressly modular: Mask R-CNN "can simultaneously predict boxes, segments, and keypoints." EX1008, 8; EX1003, ¶¶420-421.

Both architectures use shared convolutional features feeding modular heads with separate losses. A POSITA would simply attach He's conv/deconv keypoint-regression branch as another prediction head to Zhengyang's backbone and co-train with labeled keypoints, exactly how He describes training end-to-end. EX1003, ¶¶422-423. This combination is nothing more than the "predictable use of prior art according to their established functions" and thus obvious under *KSR*. *KSR*, 550 U.S. at 417. He already shows Mask R-CNN can output boxes, segments, and keypoints in parallel, confirming multi-task heads coexist naturally. A POSITA therefore would reasonably expect success. EX1003, ¶¶424-428.

7. [1G]

As noted above, Petitioner is not stating [1G] to recite a "joint loss," and regardless, Zhengyang-He renders this obvious. EX1003, ¶¶429-430

a. Zhengyang

Zhengyang teaches standard end-to-end supervised learning in which each head's output is "compared to the labeling results" and the "difference [is] reduced through the corresponding loss function … to optimize network parameters." EX1007, ¶¶[0022]-[0023]. A POSITA would understand this means losses from the heads backpropagate through shared feature maps, necessarily updating backbone weights. Zhengyang confirms "attribute classification losses … are reduced through forward and back propagation" and that per-task losses are

combined into a total weighted objective. EX1007, ¶¶[0084], [0128]; EX1003, ¶¶431-438.

b. He and MTC

He likewise trains the pose-estimation head "by minimizing cross-entropy loss" over keypoints. EX1008, 7.

MTC: For the reasons discussed for [1F], adding He's keypoint-loss term into Zhengyang's already-disclosed multi-loss objective would have been obvious and is a routine multi-task extension: same backbone, same loss-minimization, simply an additional loss head and no architectural change is required. This is the exact "combination of familiar s according to known methods to yield predictable results" that *KSR* deems obvious. *KSR*, 550 U.S. at 416; EX1003, ¶¶439-442.

8. 1[H]

Zhengyang discloses [1H]. After minimizing task-specific losses "through repeated forward and backward propagation," "the network converges … and the parameters in the network are the required final results." EX1007, ¶[0090]. These "parameters" are the trained weights of both the shared backbone and the task-specific heads obtained via loss minimization. Zhengyang states the system "select[s] the network architecture and model with the best attribute analysis results … deleting the network that does not participate… to reduce the network model size." EX1007, ¶[0012]. This shows the trained model (including backbone

+ heads) is retained after training.

Zhengyang then teaches storing and executing that trained model: "after the training… the model is deployed in a real-time monitoring system," where "the pre-trained model is used to extract feature information … and attribute recognition results are output." EX1007, ¶¶[0141]-[0144]. Zhengyang confirms the method "is stored in a computer-readable storage medium." EX1007, ¶[0146]. The stored "program" includes the trained backbone and head parameters needed for inference.

Zhengyang discloses storing, after minimizing the loss, a trained distraction-classification model comprising backbone-prediction-head parameters. EX1003, ¶¶443-449.

B. Claim 2

Zhengyang states that each distraction-classification head (smoking, phone) includes convolutional, pooling, and fully connected layers. EX1003, ¶¶450-451. This follows from: (1) identification of SE-ResNet-18 as the backbone, EX1007, ¶[0100]; (2) statements the attention networks "have the same network structure as the main network model," EX1007, ¶[0103]; and (3) per-task feature extraction, pooling, and classification. A POSITA would understand each attention head as an independent CNN sub-network trained jointly with the backbone. EX1007, ¶¶[0031]-[0035], [0100]-[0103], [0113]-[0116].

Convolutional-Layer. SE-ResNet-18 is a convolutional architecture with conv/BN/activation blocks and kernel/stride modifications. EX1007, ¶[0100]. The attention heads "have the same network structure," EX1007, ¶[0103], and perform analysis on cropped images, EX1007, ¶[0035], so a POSITA would understand each head replicates SE-ResNet-style convolution. EX1003, ¶¶452-464.

Pooling-Layer. Zhengyang states "after the four feature maps are pooled, the feature vectors are input into four attribute-recognition modules." EX1007, ¶[0101]. Pooling occurs within each head's path generating the vector input to the classifier. While the "first max-pooling layer is deleted," EX1007, ¶[0100], later pooling / stride-2 / GAP is retained consistent with SE-ResNet-18. EX1003, ¶¶465-470.

Fully Connected-Layer. Zhengyang identifies FC layers in the attention-proposal sub-network ("two fully connected layers in series," EX1007, ¶¶[0034], [0101], and FC+softmax in the final classifier, EX1007, ¶¶[0035], [0113]-[0116]. EX1003, ¶¶471-476.

Even if alleged Zhengyang does not enumerate each layer type within the heads, a POSITA would still find them obvious based on (1) Zhengyang's SE-ResNet framework and (2) standard CNN practice. Zhengyang's SE-ResNet-18 backbone, EX1007, ¶[0100], reflects the conventional convolution→pooling→fully-connected pipeline, and Zhengyang states the

attention networks "have the same network structure." EX1007, ¶[0103]. Implementing the same layer composition in each head would maintain gradient consistency, reduce training complexity, and accords with Zhengyang's teaching. By the '243 patent priority date, CNNs for driver-monitoring universally used this sequence: convolution for spatial extraction, pooling for invariance, fully connected for classification. Because the attention heads share the same input modality and feature scale as the SE-ResNet backbone, a POSITA would expect this same sequence to train stably and yield accurate classification outputs. EX1003, ¶¶477-482.

C. Claim 3

Zhengyang discloses this claim. Zhengyang teaches an object-detection prediction head composed of two sub-networks: (1) the APN, which performs bounding-box regression by predicting coordinates of an attention region; and (2) downstream attention networks, which classify the cropped region. EX1007, ¶¶[0031]-[0035], [0101]-[0103], [0113]-[0116]; EX1003, ¶¶483-484.

Bounding-Box Regression Network. The APN "comprises two fully connected layers … used to output the center-point coordinates … and the edge lengths of the attention area." EX1007, ¶[0034]; *see also* EX1007, ¶[0101]. These four outputs (x, y, width, height) are standard bounding-box parameters used in R-CNN/SSD-style regression. Zhengyang explains that these coordinates guide

cropping/enlargement, which are standard post-regression steps. Thus, the APN provides bounding-box localization within Zhengyang's multi-task architecture. EX1003, ¶¶485-488.

Object-Class Prediction Network. The localized region is then processed by attention networks that "re-judge the attributes." EX1007, ¶[0102]; *see also* EX1007, ¶[0033]. This is categorical classification. Zhengyang states these attention models "have the same network structure as the main network model," i.e., SE-ResNet-18 with soft-max classification output. EX1007, ¶¶[0103], [0100]. Zhengyang identifies binary class outputs, EX1007, ¶[0028], and soft-max probability computation, EX1007, ¶¶[0113]-[0116], confirming that the attention networks output class probabilities, i.e., object-class predictions. EX1003, ¶¶489-493.

Zhengyang discloses an integrated detection head: the APN predicts where (bounding-box parameters) and the attention networks predict what (class label). EX1003, ¶494.

D. Claim 4

Zhengyang discloses this claim. Zhengyang discloses both the bounding-box regression network and the object-class prediction network are deep neural networks with multiple hidden layers comprising convolution, batch-norm, and activation. Zhengyang defines the model as SE-ResNet-18 and states its sub-

networks share the same structure, including the attention-proposal and attention-classification modules. EX1007, ¶¶[0033]-[0035], [0100]-[0103]. Both networks adopt SE-ResNet residual blocks consisting of repeated convolutional-Batch_normalization-activation layers. EX1003, ¶¶495-500.

SE-ResNet-18 layer template. Zhengyang states "the multi-task network structure used here is SE-ResNet-18 … BN layer, PReLU, convolution … ." EX1007, ¶[0100]. This as the SE-ResNet block pattern, convolutional, Batch_Normalization, activation repeated across depth. EX1003, ¶¶501-504.

Bounding-box regression network (APN). Although Zhengyang states the APN "mainly comprises two fully connected layers," EX1007, ¶[0034], a POSITA would understand those to be *final* regression layers. Consistent with region-proposal practice, earlier feature-processing layers necessarily use SE-ResNet-style convolutional-Batch_normalization-activation blocks. EX1007, ¶¶[0034], [0100]; EX1003, ¶¶505-507.

Object-class prediction network (attention networks). Zhengyang states the attention networks "have the same network structure as the main network model." EX1007, ¶[0103]. Because the main network is SE-ResNet-18, EX1007, ¶[0100], each attention network likewise repeats convolutional-Batch_normalization-activation layers. EX1007, ¶¶[0100]-[0103]; EX1003, ¶¶508-509.

Thus, Zhengyang's explicit "BN layer, PReLU, convolution" disclosure,

EX1007, ¶[0100], and its statement that the attention networks share the same structure, EX1007, ¶[0103], would lead a POSITA to understand that both the APN and the attention networks employ repeated convolutional-Batch_normalization-activation hidden layers, consistent with SE-ResNet and standard Faster-R-CNN-style practice. EX1007, ¶¶[0033], [0100]-[0103]; EX1003, ¶¶510-517.

E. Claim 5

Zhengyang-He renders this claim obvious. EX1003, ¶518.

1. Zhengyang

As established in Claim 1, Zhengyang's multi-task SE-ResNet-18 framework is designed for multiple task heads and is amenable to adding a pose-estimation head. Zhengyang provides the shared backbone, multi-head context, and explicit layer template of "BN layer, PReLU, convolution." EX1007, ¶[0100]; EX1003, ¶¶519-520.

2. He

Even if Zhengyang is alleged deficient, He provides a concrete pose-estimation head that operates as a deep convolutional sub-network with multiple hidden layers and a final convolutional output that produces keypoint heatmaps. EX1008, 7-8; EX1003, ¶¶521-525.

Hidden-layer composition in a pose-estimation head. He teaches a keypoint head that "consists of a stack of eight 3×3 512-d convolution layers." EX1008, 8. That is a plurality of hidden layers. In SE-ResNet based heads, each convolution layer is paired with a batch-normalization layer and an activation layer.

MTC: Zhengyang states SE-ResNet-18 uses batch-normalization, PReLU activation, and convolution. EX1007, ¶[0100]. Zhengyang also states task heads "have the same network structure as the main network model." EX1007, ¶[0103]. Therefore, a POSITA would implement each hidden layer in the added pose head as a convolution layer, a batch-normalization layer, and an activation layer in order to maintain consistency with the SE-ResNet-18 template and to maintain stable feature scaling across all heads. EX1007, ¶¶[0100], [0103]; EX1003, ¶526.

Output layer of the pose-estimation head. He teaches that the keypoint head produces K keypoint masks and details a final convolutional stage that generates the heatmap outputs, with upsampling via deconvolution and bilinear scaling to arrive at the target resolution. EX1008, 7-8. A POSITA would recognize this as an output layer that comprises a convolutional layer, satisfying the claim's output-layer requirement. EX1003, ¶¶527-529.

MTC: A POSITA would attach He's keypoint head as an additional branch to Zhengyang's shared SE-ResNet-18 feature maps and implement the hidden layers using the same convolution-BN-activation block pattern already specified for

Zhengyang's network. EX1007, ¶[0100]; EX1008, 8. **Further,** Zhengyang seeks to "comprehensively reflect the driver's driving state and improve the accuracy of driver state recognition." EX1007, ¶[0046]. Adding a pose-estimation head yields keypoint heatmaps for head, shoulders, arms, and related joints, which directly improves assessment of posture and distraction while coexisting with other task heads, as demonstrated in multi-task frameworks. EX1008, 8. **Finally,** noth references use convolutional backbones with task-specific heads, and the layer types are standard modules. He reports successful joint operation of detection, segmentation, and keypoints in one model, indicating that adding a pose branch does not disrupt other tasks. *Id.* Given Zhengyang's explicit "BN layer, PReLU, convolution" template and shared-structure statement, a POSITA would expect straightforward integration and stable training. EX1007, ¶¶[0100], [0103]; EX1003, ¶¶530-534.

F. Claim 6

Petitioner is not reading Claim 6 to recite a dedicated "distraction-classification prediction head," and regardless, Zhengyang-He renders obvious claim 6.

Zhengyang teaches inputting driver-facing images into a trained multi-task SE-ResNet-18 network and outputting binary attribute results including whether there is smoking and whether a phone call is being made. EX1007, ¶¶[0020]-

[0023], [0028], [0031], [0046], [0061]-[0062], [0100]-[0103], [0113]-[0116]. Zhengyang states each head produces "two categories" (e.g., smoking vs. not smoking) using a soft-max output. EX1007, ¶¶[0028], [0113]-[0116]. These binary outputs are the claimed distraction-classification "tag." And even if it were argued that Zhengyang does not expressly disclose receiving a "tag," it would have been obvious in view of He.

He confirms that this is the standard output produced by each CNN classification branch: He discloses heads that output both "class label" and associated offsets, selected via soft-max. EX1008, 3, 8. He explains that the model uses the mask associated with the predicted class selected by the classification branch. EX1008, 8. This is the conventional sequence of a convolutional head followed by a fully connected layer and soft-max that yields a discrete class label. That mechanism corresponds to the "receiving a distraction classification tag" step, because both produce categorical identifiers derived from image data.

MTC. He's explicit soft-max output therefore confirms the label-generation step implicit in Zhengyang's attribute-recognition results. EX1003, ¶¶535-546. A POSITA would integrate He's explicit soft-max class-label output into Zhengyang's attention-head architecture because both references use the same shared-backbone, multi-head CNN framework with convolution, BN, activation, and fully connected layers. EX1007, ¶¶[0100]-[0103]; EX1008, 3. Zhengyang

seeks to "comprehensively reflect the driver's driving state and improve the accuracy of driver state recognition." EX1007, ¶[0046]. Generating categorical labels (e.g., smoking vs. not-smoking) directly advances that objective, and He confirms such class-label heads are standard in modern CNNs and operate in real-time within multi-head models. EX1008, 3, 8. A POSITA therefore would have had a reasonable expectation of success. EX1003, ¶¶547-553.

G. Claim 9

Zhengyang-He renders claim obvious 9. EX1003, ¶554

1. Zhengyang

Zhengyang teaches a multi-task CNN that classifies driver images into categorical distraction states, such as smoking or phone use, to improve driving safety. EX1007, ¶¶[0002], [0021]. These outputs are binary yes or no for each behavior, which are categorical labels derived from the feature vector of each image, consistent with standard soft max classification. EX1007, ¶[0028]. Smoking and phone use correspond to distracted states, and the absence of those behaviors corresponds to undistracted. Zhengyang identifies smoking recognition and phone call recognition heads, which map directly to the claimed mobile usage and other distracted categories. EX1007, ¶[0008]; EX1003, ¶¶555-559.

2. He

He confirms that CNN classification heads produce class labels. Mask R CNN outputs a class label selected by the classification branch. EX1008, 3, 8.

3. MTC

A POSITA would therefore incorporate He's explicit labeling structure into Zhengyang's prediction heads to formalize the categorical outputs already described in Zhengyang. This is obvious because both references use shared backbones and task specific heads that feed fully connected soft max decision layers. EX1003, ¶560. It would have been straightforward to combine these teachings because He's class label format merely clarifies the binary outputs generated by Zhengyang's heads, and directly supports Zhengyang's stated goal to improve driving safety by identifying distraction states. EX1007, ¶[0002]. A POSITA would reasonably expect success because by the priority date, CNN driver monitoring systems were already labeling driver states such as attentive, phone use, or other distraction. EX1003, ¶¶561-565.

H. Claim 10

Zhengyang attention proposal sub network discloses claim 10. The APN performs object localization by outputting spatial parameters indicating where a target object appears in the driver image. Zhengyang states that it comprises two fully connected layers that output the center point coordinates and edge lengths of the attention area for cropping and enlargement of the original input image. EX1007, ¶[0034]. The "input image" is represented as a numerical feature vector produced by the convolutional backbone and that the APN outputs the coordinates

that identify the location of the object within that vector.

Zhengyang further describes iterative refinement in which the APN's predicted region is cropped and reanalyzed for improved localization accuracy. EX1007, ¶¶[0101]-[0102]. This sis the standard region proposal and bounding box refinement loop used in conventional object detection.

Accordingly, the APN detects the location of an object such as a phone or cigarette within the input feature vector. EX1003, ¶¶566-572.

I. Claim 11

Zhengyang-He renders obvious claim 11. EX1003, ¶573

Zhengyang converts each driver image into a numerical feature representation through convolutional processing, which is the input vector on which each head operates. EX1003, ¶574.

He teaches a modular pose estimation head receives shared convolutional features and outputs keypoint heatmaps representing the orientation and location of human body parts. He explains the model is extended to pose estimation and describes a keypoint head with stacked convolutional layers followed by deconvolution and bilinear upscaling to produce spatial heatmaps. EX1008, 7-8. These heatmaps represent the orientation of body parts derived directly from the input vector. EX1003, ¶¶575-577.

MTC: Zhengyang already generates localized attention regions focused on the driver's head, hands, and face, which could be used as inputs to a pose estimation head of the type disclosed by He. Integrating this head into Zhengyang's multi task framework would require minimal modification because both references use shared backbones and modular heads. EX1003, ¶¶578-580.

A POSITA would be motivated to add He's pose estimation head to improve Zhengyang's ability to assess distraction, since orientation cues such as head turn or arm position are relevant to attentiveness. A POSITA would expect success due to the architectural compatibility, and the combined system would detect orientation of body parts in the input vector as required by claim 11. EX1003, ¶¶581-585.

J. Claim 12

Zhengyang discloses claim 12. EX1003, ¶586. As discussed for claim 1, Zhengyang teaches storing a trained distraction-classification model comprising backbone and head parameters for inference. Storing that model in the dash mounted camera is a routine implementation. EX1003, ¶¶587-588.

Zhengyang discloses an in-vehicle camera device with onboard processing and memory that acquires the face image and inputs it to the network, confirming a camera module containing sensor, processor, and memory. EX1007, ¶¶[0014], [0045]. Zhengyang also prunes the model to reduce size and forward transmission

time, which indicates local inference, and therefore storage within device memory. EX1007, ¶¶[0038]-[0039]; EX1003, ¶¶589-591. Based on these teachings, Zhengyang's trained model is stored within the dash-mounted camera to minimize latency, reduce wiring complexity, and enable immediate driver-state detection.

K. Claims 13-18 and 20-24

Zhengyang itself discloses that the same multi-task driver state recognition method is embodied as a CRM storing executable instructions, and as a processor-and-memory-based device executing those instructions. EX1007, ¶¶[0044]-[0045], [0137]-[0142], [0143]-[0147]. Thus, for the same reasons discussed above for claims 1-7 and 9-12, a POSITA would have found storing program instructions on a medium (claims 13-18) or executing those same instructions on a processor in a device (claims 20-24) obvious in view of Zhengyang-He. Accordingly, claims 13-18 and 20-24 are obvious for the same reasons as claims 1-7 and 9-12. EX1003, ¶¶592-600.

VIII. GROUND 4: ZHENGYANG IN VIEW OF HE AND FURTHER IN VIEW OF CHAO RENDERS OBVIOUS CLAIMS 7-8, 19, 25-33.

A. Claims 7, 8, 19 and 25

Zhengyang in view of He and further in view of Chao renders obvious claims 7, 8, 19 and 25. Claims 19 and 25 merely restate the same functional recitations of claims 7 and 8, but in CRM and device form. EX1003, ¶¶601-603

1. Zhengyang-He

A POSITA would understand that Zhengyang-He discloses, teach or suggest claims 7, 8, 19 and 25.

Zhengyang produces a distraction-classification tag (for example "smoking" or "phone call") and teaches these categorical outputs are generated in an intelligent vehicle environment "so as to improve driving safety." EX1007, ¶[0002]. Zhengyang explicitly identifies the distraction states "whether there is smoking, and whether a phone call is being made," EX1007, ¶¶[0021], [0028], and describes these outputs "reflect the driver's driving state and improve the accuracy of driver state recognition," EX1007, ¶[0046]. A POSITA would not interpret these categorical tags as inert values. In a safety motivated IVMS used "to improve driving safety," such categorical tags function as decision tokens that drive the next operation in the pipeline. EX1003, ¶¶604-606.

Zhengyang also discloses an integrated in vehicle camera device with a processor, memory, and "an output apparatus with a display and speaker." EX1007, ¶[0152]; *see also* EX1007, ¶¶[0149]-[0153]. This configuration is a driver monitoring module that performs inference and has the output pathways needed to consume and act on the distraction tag. EX1003, ¶¶607-609.

He independently reinforces that the predicted class is the control variable that determines what happens next. He states that the network outputs "a class

label and a bounding box offset" and that the system "only use[s] the k-th mask where k is the predicted class." EX1008, 3, 8. This is standard gating logic. Under the breadth of claim 7, this gating behavior itself is an action based on the distraction-classification tag.

For that reason, Zhengyang-He renders obvious claim 7. EX1003, ¶¶610-615.

If it is alleged claim 7 recites an outward user facing action such as visual, audible, or network transmission output (i.e., claim 8), that implementation would have been obvious further in view of Chao, which explicitly teaches the exact downstream actuation contemplated by claims 7, 8, 19, and 25.

2. Chao

Chao teaches a driver-monitoring CNN that outputs a class label and then uses that determined state to "trigger[] vehicular actions based on the determined state." EX1006, ¶¶[0022], [0068], [0095], [0129]. Those actions include notifications "via the UIs," which "may include a visual notification, a sound notification, or a vibration notification." EX1006, ¶¶[0095], [0117], [0168]. A visual notification is updating a display, a sound notification is an audible alarm, and that the determined state is delivered to a separate processor. EX1006, ¶¶[0032], [0075], [0078].

Accordingly, Chao provides the mechanism that maps the classification output to display updates, alarms, or remote transmission, which are the actions recited in claim 8. EX1003, ¶¶616-620.

3. MTC

A POSITA would view Zhengyang-He-Chao as three adjacent stages of the same pipeline: Zhengyang generates the categorical distraction output, He shows that a class ID is the control signal that determines what operation executes next, and Chao consumes the class/state to drive HMI or vehicle actuation. The integration is mechanically trivial: the discrete class value or softmax output produced by the Zhengyang task head would be passed directly into the notification or actuation module described in Chao, with no new network layers, no feature-map changes, and no data-type adaptation. All three references assume the same modality (driver-facing images), the same architecture family (multi-head CNN), and the same deployment surface (in-vehicle runtime with display/speaker/control channels). The technical alignment across the three references means a POSITA would expect success on first implementation, because this is the standard deployment pattern: perception, classification, then actuation *KSR*, 550 U.S. at 416-417. Claims 7, 8, 19 and 25 therefore add no features beyond routine label-based control that a POSITA would have expected to apply. EX1003, ¶¶621-625.

B. Claims 26-33

Claims 26-33 recite no new features, rather they merely restate the same limitations of claims 1-12 in the context of a dash-mounted in-vehicle inference device. Independent claims 26 and 30 are the method/device counterparts, and claims 27-29 and 31-33 track the same refinements that appear in claims 9-11.

Claim 26 mirrors the method limitations of claims 1, 6-8, and 12, but placed in a dash-mounted camera device. For the same reasons already discussed for those claims, storing the trained model in the dash-mounted unit, capturing in-cabin images, inputting the images into the model, receiving the distraction tag, and performing an action based on the tag would have been obvious in view of Zhengyang-He-Chao. Zhengyang describes the same in-vehicle capture→inference pipeline, EX1007, ¶¶[0060], [0109], [0145]; He confirms multi-head architectures trained end-to-end; and Chao expressly teaches triggering actions based on driver-state classification, including "visual," "sound," or remote notifications. EX1006, ¶¶[0095], [0168]; EX1003 ¶¶626-631.

Claims 27-29 add no features beyond claim 26 and correspond to the same refinements addressed in claims 9-11 (e.g., classification labels, bounding-box location, and pose orientation) using image data instead of an input vector, which would have been obvious for the same reasons in view of Zhengyang-He-Chao as the input vector is generated based on image data of drivers. EX1003, ¶632.

Claim 30 simply recasts claim 26 as a processor-and-memory device. Zhengyang already discloses the same trained backbone+heads model embodied as executable instructions stored in memory and executed by a processor within an in-vehicle unit. EX1007, ¶¶[0044]-[0045], [0137]-[0146]. Claims 31-33 add the same dependent refinements as claims 27-29, again with no new substance. EX1003, ¶¶633-635.

Accordingly, claims 26-33 merely restate the same features in CRM/device form. A POSITA would have found each obvious for the reasons already articulated in view of Zhengyang-He-Chao.

IX. NO OBJECTIVE INDICIA OF PATENTABILITY

Petitioner is unaware of any objective indicia of non-obviousness. If alleged in the POPR, Petitioner requests a response. EX1003, ¶¶636-639.

X. MANDATORY NOTICES

A. Real Party in Interest

Petitioner Samsara Inc. ("Samsara") is the only real party-in-interest.

B. Related Matters

Motive asserted the '243 patent in the Northern District of California (Case No. 3-24-cv-00902) in an amended complaint filed on July 9, 2025. EX1037. In the same amended complaint, Motive also asserted U.S. Patent No. 12,136,276 (the "'276 patent"). The '276 patent is a child of the '580 patent. An IPR trial (IPR2025-00574) was instituted against the '580 patent on August 27, 2025. An IPR was filed against the '276 patent on October 17, 2025 (IPR2026-00034).

A. Lead and Backup Counsel

Lead Counsel	**Back-up Counsel**
Jason D. Eisenberg (Reg. No. 43,447) jasone-PTAB@sternekessler.com STERNE, KESSLER, GOLDSTEIN & FOX PLLC 1101 K Street, NW, 10th Floor Washington, DC 20005 202.371.2600 (reception) 202.371.2540 (facsimile)	Lestin Kenton (Reg. No. 72,314) lkenton-PTAB@sternekessler.com Bryn Mendenhall (Reg. No. 82,263) bmendenhall-PTAB@sternekessler.com STERNE, KESSLER, GOLDSTEIN & FOX PLLC 1101 K Street, NW, 10th Floor Washington, DC 20005 202.371.2600 (reception) 202.371.2540 (facsimile)

Please address all correspondence to lead and backup counsel at the email addresses shown above and PTAB@sternekessler.com. Samsara consents to electronic service by email.

XI. GROUNDS FOR STANDING

Samsara certifies the '243 patent is available for IPR and Samsara is not barred or estopped from requesting IPR of the '243 patent challenging the patent claims on the Grounds identified in this Petition.

XII. CONCLUSION

The Board should institute *inter partes* review and cancel all challenged claims.

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg, Reg. No. 43,447
Lestin Kenton, Reg. No. 72,314
Bryn Mendenhall, Reg. No. 82,263
Counsel for Petitioner Samsara Inc.

Date: November 14, 2025

1101 K Street NW, 10th Floor
Washington, DC 20005
(202) 371-2600

Claims Appendix – U.S. Patent No. 12,062,243

Claim 1

Element	
[1P]	A method comprising
[1A]	inputting labeled examples into a neural network, the neural network comprising:
[1B]	a backbone network, and
[1C]	one or more prediction heads coupled to the backbone network, the one or more prediction heads comprising one or more of
[1D]	a distraction classification prediction head
[1E]	object detection prediction head
[1F]	and a pose estimation prediction head
[1G]	minimizing a loss based on outputs of the one or more prediction heads causing a change in parameters of the backbone network; and
[1.H]	storing a distraction classification model after minimizing the loss, the distraction classification model comprising the parameters of the backbone network and parameters of at least one of the one or more prediction heads.

Claim 2

Element	
[2]	The method of claim 1, wherein the distraction classification prediction head includes a convolutional layer, pooling layer, and fully connected layer.

Claim 3

Element	
[3]	The method of claim 1, wherein the object detection prediction head comprises a bounding box regression network and an object class prediction network, wherein the bounding box regression network outputs coordinates of a bounding box enclosing a detected object and the object class prediction network outputs a class corresponding to the detected object.

Claim 4

Element	
[4]	The method of claim 3, wherein each of the bounding box regression network and the object class prediction network comprise deep neural networks having a plurality of hidden layers, each hidden layer in the plurality of hidden layers comprising a convolutional layer, a batch normalization layer, and a batch activation layer.

Claim 5

Element	
[5]	The method of claim 1, wherein the pose estimation prediction head includes a plurality of hidden layers and an output layer, each of the plurality of hidden layers comprising a convolutional layer, a batch normalization layer, and an activation layer, and the output layer comprising a convolutional layer.

Claim 6

Element	
[6]	The method of claim 1, further comprising inputting image data into the distraction classification model and receiving a distraction classification tag from the distraction classification model based on the image data.

Claim 7

Element	
[7]	The method of claim 6, further comprising performing at least one action based on the distraction classification tag.

Claim 8

Element	
[8]	The method of claim 7, the at least one action comprising one or more of updating a display of a computing device, playing an audible alarm, transmitting or a prediction to a remote server.

Claim 9

Element	
[9]	The method of claim 1, wherein the distraction classification prediction head classifies an input vector with a label, the label comprising one or more of an undistracted, mobile usage, or other distracted label.

Claim 10

Element	
[10]	The method of claim 1, wherein the object detection prediction head detects a location of an object within an input vector.

Claim 11

Element	
[11]	The method of claim 1, wherein the pose estimation prediction head detects an orientation of one or more body parts of a user depicted in an input vector.

Claim 12

Element	
[12]	The method of claim 1, wherein storing a distraction classification model comprises storing the distraction classification model in a dash-mounted camera device.

Claim 13

Element	
[13P]	A non-transitory computer-readable storage medium for tangibly storing computer program instructions capable of being executed by a computer processor, the computer program instructions defining steps of:
[13A]	inputting labeled examples into a neural network, the neural network comprising:
[13B]	a backbone network, and one or more prediction heads,
[13C]	the one or more prediction heads selected from a group consisting of a distraction classification prediction head, object detection prediction head, and a pose estimation prediction head;
[13D]	a distraction classification prediction head,
[13E]	object detection prediction head,
[13F]	and a pose estimation prediction head;
[13G]	minimizing a loss based on outputs of the one or more prediction heads causing a change in parameters of the backbone network; and
[13H]	storing a distraction classification model after minimizing the loss, the distraction classification model comprising the parameters of the backbone network and parameters of at least one of the one or more prediction heads.

Claim 14

Element	
[14]	The non-transitory computer-readable storage medium of claim 13, wherein the distraction classification prediction head includes a convolutional layer, pooling layer, and fully connected layer.

Claim 15

Element	
[15]	The non-transitory computer-readable storage medium of claim 13, wherein the object detection prediction head comprises a bounding box regression network and an object class prediction network, wherein the bounding box regression network outputs coordinates of a bounding box enclosing a detected object and the object class prediction network outputs a class corresponding to the detected object.

Claim 16

Element	
[16]	The non-transitory computer-readable storage medium of claim 15, wherein each of the bounding box regression network and the object class prediction network comprise deep neural networks having a plurality of hidden layers, each hidden layer in the plurality of hidden layers comprising a convolutional layer, a batch normalization layer, and a batch activation layer.

Claim 17

Element	
[17]	The non-transitory computer-readable storage medium of claim 13, wherein the pose estimation prediction head includes a plurality of hidden layers and an output layer, each of the plurality of hidden layers comprising a convolutional layer, a batch normalization layer, and an activation layer, and the output layer comprising a convolutional layer.

Claim 18

Element	
[18]	The non-transitory computer-readable storage medium of claim 13, the steps further comprising inputting image data into the distraction classification model and receiving a distraction classification tag from the distraction classification model based on the image data.

Claim 19

Element	
[19]	The non-transitory computer-readable storage medium of claim 18, the steps further comprising performing at least one action based on the distraction classification tag.

Claim 20

Element	
[20P]	A device comprising:
[20A]	a processor; and
[20B]	a storage medium for tangibly storing thereon program logic for execution by the processor, the program logic comprising:
[20C]	logic, executed by the processor, for inputting labeled examples into a neural network, the neural network comprising: a backbone network, and one or more prediction heads, the one or more prediction heads selected from a group consisting of a distraction classification prediction head, object detection prediction head, and a pose estimation prediction head,
[20D]	logic, executed by the processor, for minimizing a loss based on outputs of the one or more prediction heads causing a change in parameters of the backbone network, and
[20E]	logic, executed by the processor, for storing a distraction classification model after minimizing the loss, the distraction classification model comprising the parameters of the backbone network and parameters of at least one of the one or more prediction heads.

Claim 21

Element	
[21]	The device of claim 20, wherein the object detection prediction head comprises a bounding box regression network and an object class prediction network, wherein the bounding box regression network outputs coordinates of a bounding box enclosing a detected object and the object class prediction network outputs a class corresponding to the detected object.

Claim 22

Element	
[22]	The device of claim 21, wherein each of the bounding box regression network and the object class prediction network comprise deep neural networks having a plurality of hidden layers, each hidden layer in the plurality of hidden layers comprising a convolutional layer, a batch normalization layer, and a batch activation layer.

Claim 23

Element	
[23]	The device of claim 20, wherein the pose estimation prediction head includes a plurality of hidden layers and an output layer, each of the plurality of hidden layers comprising a convolutional layer, a batch normalization layer, and an activation layer, and the output layer comprising a convolutional layer.

Claim 24

Element	
[24]	The device of claim 20, the program logic further comprising logic, executed by the processor, for comprising inputting image data into the distraction classification model and receiving a distraction classification tag from the distraction classification model based on the image data.

Claim 25

Element	
[25]	The device of claim 24, the program logic further comprising logic, executed by the processor, for performing at least one action based on the distraction classification tag.

Claim 26

Element	
[26.P]	A method comprising:
[26A]	storing a distraction classification model in a dash-mounted camera device, the distraction classification model comprising:
[26B]	a backbone network and one or more prediction heads coupled to the backbone network, the one or more prediction heads comprising one or more of a distraction classification prediction head, object detection prediction head, and a pose estimation prediction head, wherein the distraction classification model comprises a model trained by inputting labeled examples into the distraction classification model and minimizing a loss based on outputs of the one or more prediction heads causing a change in parameters of the backbone network;
[26C]	capturing an image of an interior of a vehicle using a sensor in the dash-mounted camera device;
[26D]	inputting the image into the distraction classification model;
[26E]	receiving a distraction classification tag output by the distraction classification model; and
[26F]	performing at least one action by the dash-mounted camera device based on the distraction classification tag, the at least one action comprising one or more of updating a display of a computing device, playing an audible alarm, transmitting or a prediction to a remote server.

Claim 27

Element	
[27]	The method of claim 26, wherein the distraction classification prediction head classifies the image with a label, the label comprising one or more of an undistracted, mobile usage, or other distracted label.

Claim 28

Element	
[28]	The method of claim 26, wherein the object detection prediction head detects a location of an object within the image.

Claim 29

Element	
[29]	The method of claim 26, wherein the pose estimation prediction head detects an orientation of one or more body parts of a user depicted in the image.

Claim 30

Element	
[30P]	A dash-mounted camera device comprising:
[30A]	a memory configured to:
[30B]	store a distraction classification model, the distraction classification model comprising:
[30C]	a backbone network and one or more prediction heads coupled to the backbone network, the one or more prediction heads comprising one or more of a distraction classification prediction head, object detection prediction head, and a pose estimation prediction head, wherein the distraction classification model comprises a model trained by inputting labeled examples into the distraction classification model and minimizing a loss based on outputs of the one or more prediction heads causing a change in parameters of the backbone network;
[30D]	a sensor oriented toward an interior of a vehicle; and
[30E]	a processor configured to:
[30F]	capture an image of an interior of a vehicle;
[30G]	input the image into the distraction classification model;
[30H]	receive a distraction classification tag output by the distraction classification model; and

[30I]	perform at least one action by the dash-mounted camera device based on the distraction classification tag, the at least one action comprising one or more of updating a display of a computing device, playing an audible alarm, transmitting or a prediction to a remote server.

Claim 31

Element	
[31]	The dash-mounted camera device of claim 30, wherein the distraction classification prediction head classifies the image with a label, the label comprising one or more of an undistracted, mobile usage, or other distracted label.

Claim 32

Element	
[32]	The dash-mounted camera device of claim 30, wherein the object detection prediction head detects a location of an object within the image.

Claim 33

Element	
[33]	The dash-mounted camera device of claim 30, wherein the pose estimation prediction head detects an orientation of one or more body parts of a user depicted in the image.

Case IPR2026-00108
U.S. Patent No. 12,062,243

CERTIFICATE OF COMPLIANCE WITH TYPE-VOLUME LIMITATION, TYPEFACE REQUIREMENTS, AND TYPE STYLE REQUIREMENTS

1. This Petition complies with the type-volume limitation of 14,000 words, comprising 13,639 words, excluding the parts exempted by 37 C.F.R. § 42.24(a).

2. This Petition complies with the general format requirements of 37 C.F.R. § 42.6(a) and has been prepared using Microsoft® Word in 14-point Times New Roman.

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg
Registration No. 43,447
Attorney for Petitioner Samsara Inc.

Date: November 14, 2025

1101 K Street NW, 10th Floor
Washington, DC 20005
(202) 371-2600

CERTIFICATION OF SERVICE (37 C.F.R. §§ 42.6(e), 42.105(a))

The undersigned hereby certifies that the above-captioned **PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 12,062,243** and all supporting exhibits were served via FedEx Express® on November 14, 2025, in their entireties on the following:

GREENBERG TRAURIG (NY)
One Vanderbilt Avenue
New York, NY 10017
Correspondence Address for U.S. Patent No. 12,062,243

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg
Registration No. 43,447
Attorney for Petitioner Samsara Inc.

Date: November 14, 2025

1101 K Street NW, 10th Floor
Washington, DC 20005
(202) 371-2600

EXHIBIT C


UNITED STATES
PATENT AND TRADEMARK OFFICE
uspto

PTAB Trial Statistics

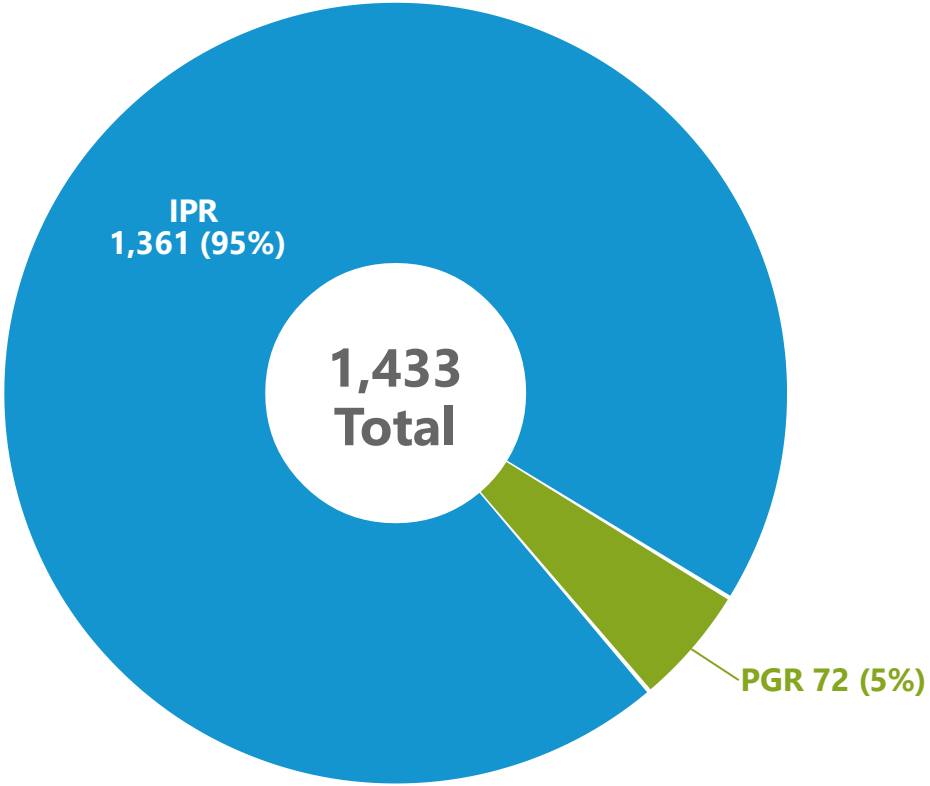
2025 End of Year Outcome Roundup

Patent Trial and Appeal Board

Fiscal Year 2025

Petitions filed by trial type

FY 25: Oct. 1, 2024 to Sep. 30, 2025

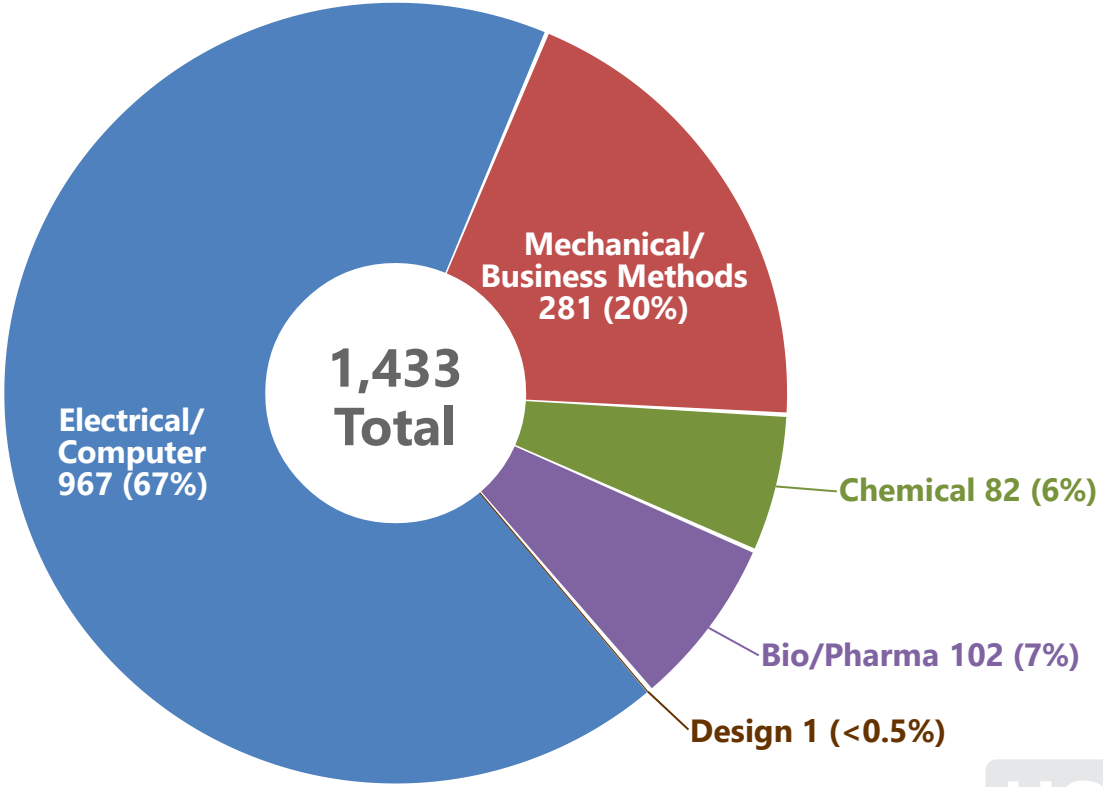


Trial types include Inter Partes Review (IPR) and Post Grant Review (PGR).

Motive Technologies Exhibit 2003, IPR2026-00034
Page 228

Petitions filed by technology

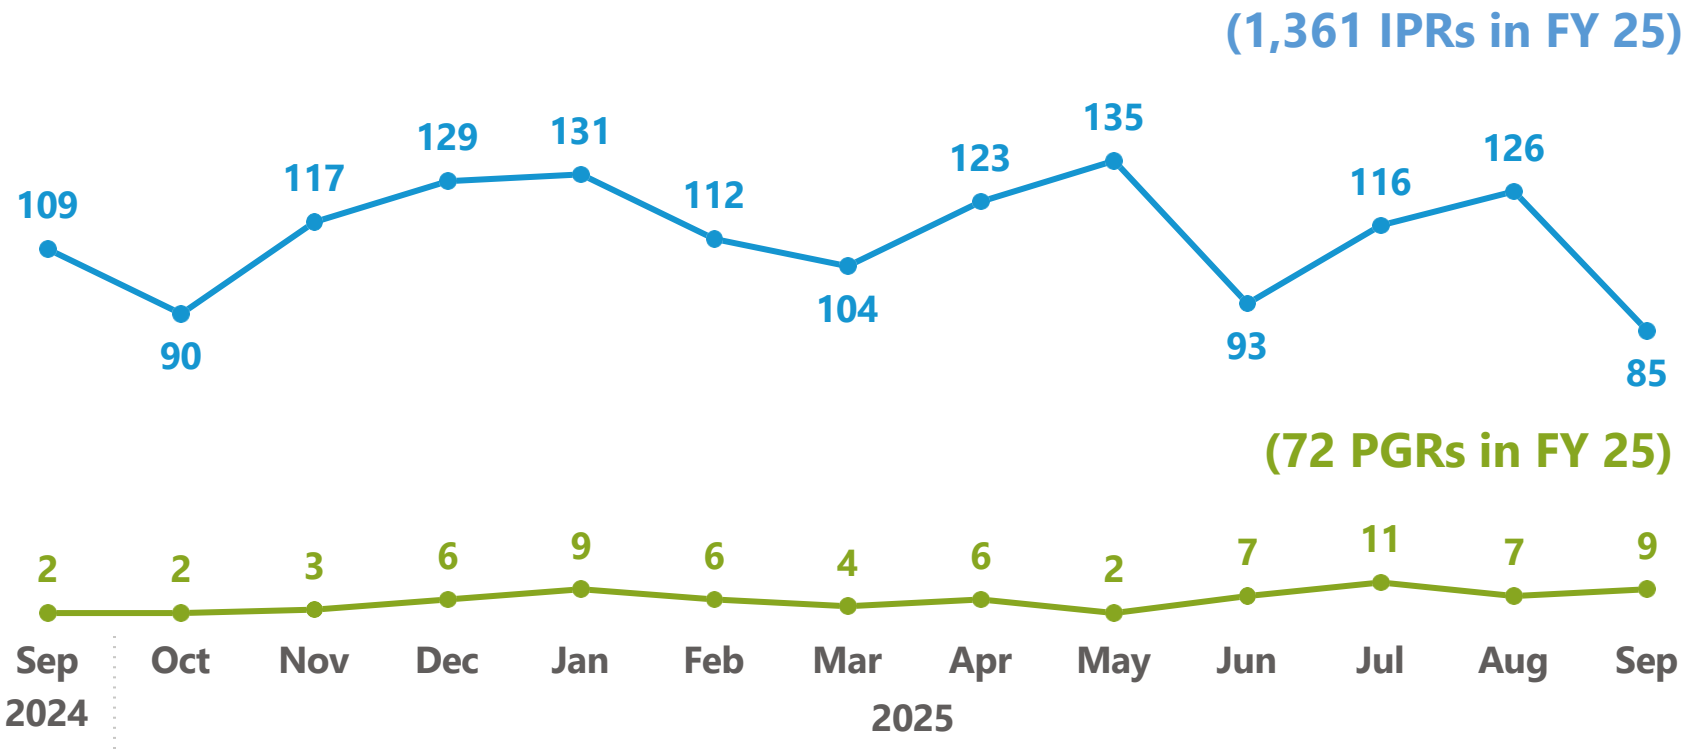
FY 25: Oct. 1, 2024 to Sep. 30, 2025



Motive Technologies Exhibit 2003, IPR2026-00034
Page 229

Petitions filed by month

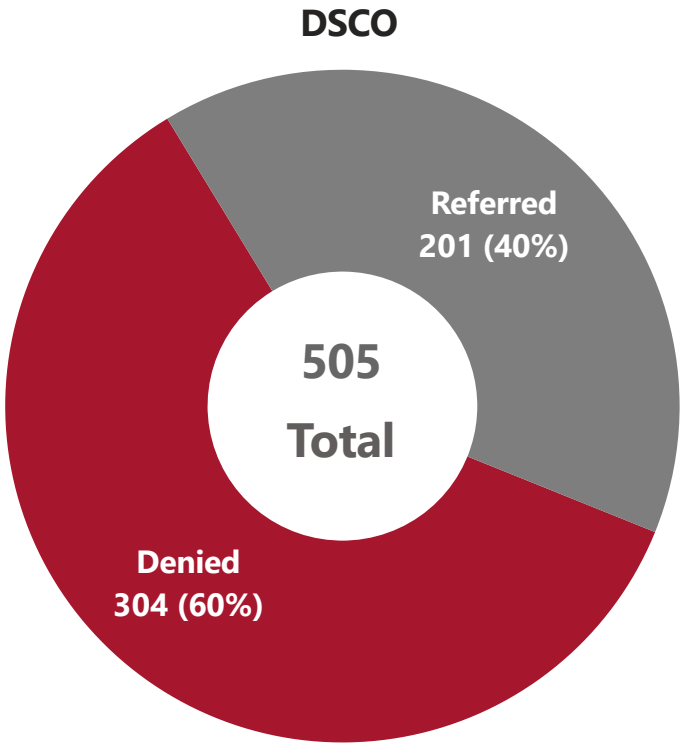
FY 24 to FY 25: Sep. 1, 2024 to Sep. 30, 2025

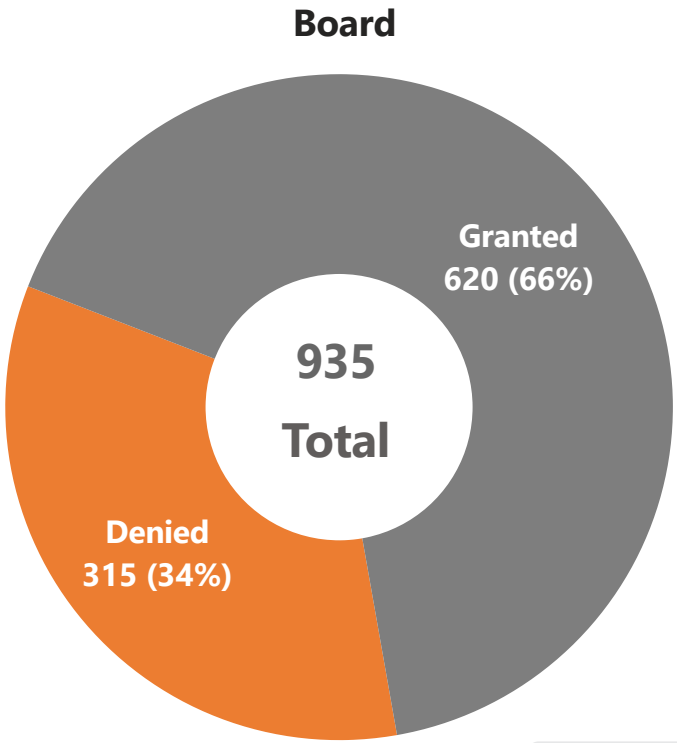


Motive Technologies Exhibit 2003, IPR2026-00034
Page 230

Disposition rates by decision type

FY 25: Oct. 1, 2024 to Sep. 30, 2025

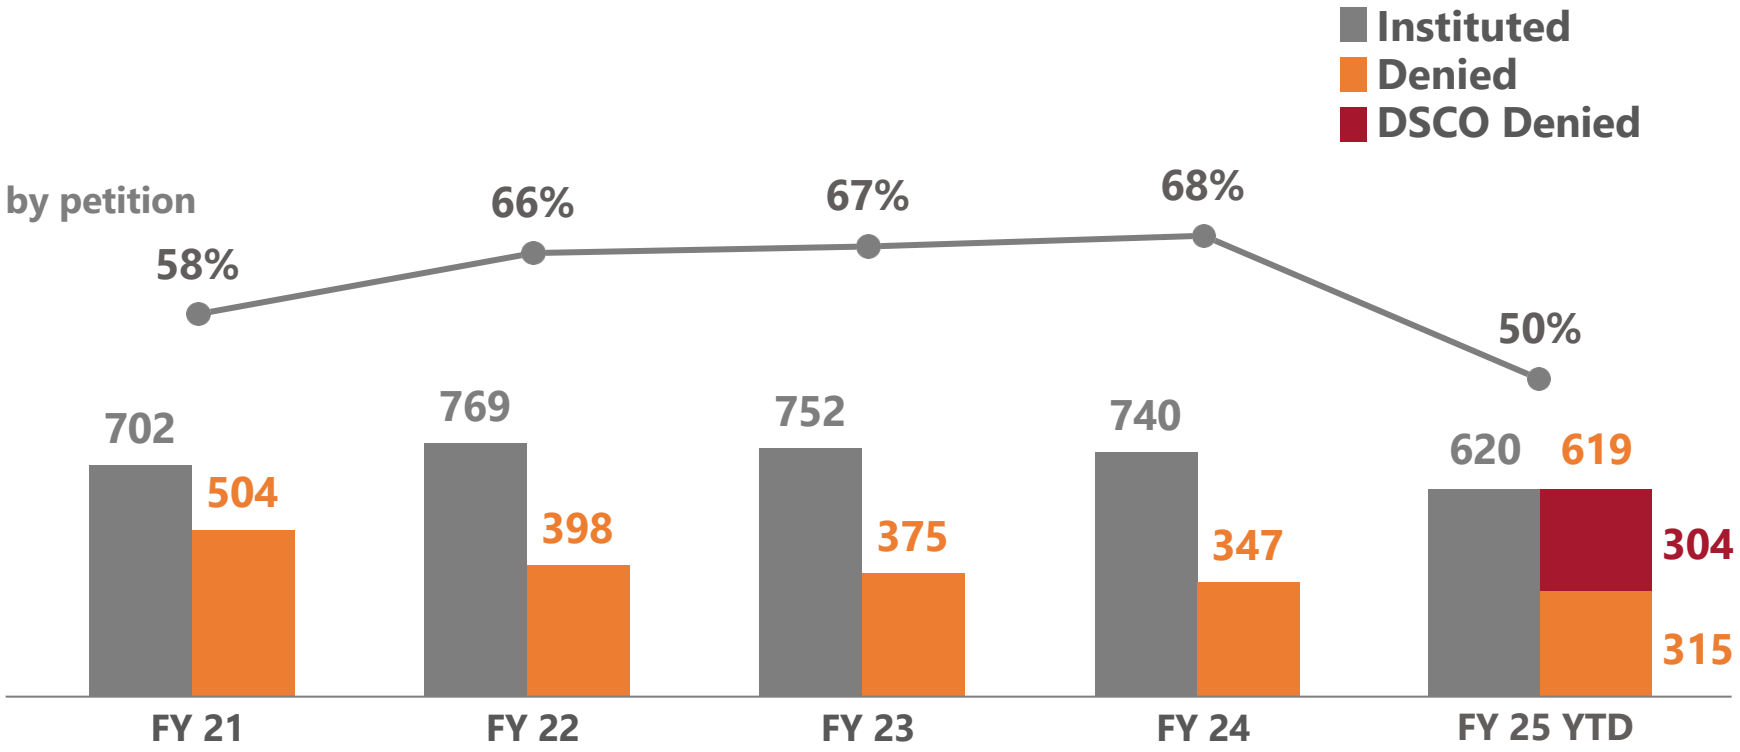




"DSCO" is "Director Discretionary Considerations"

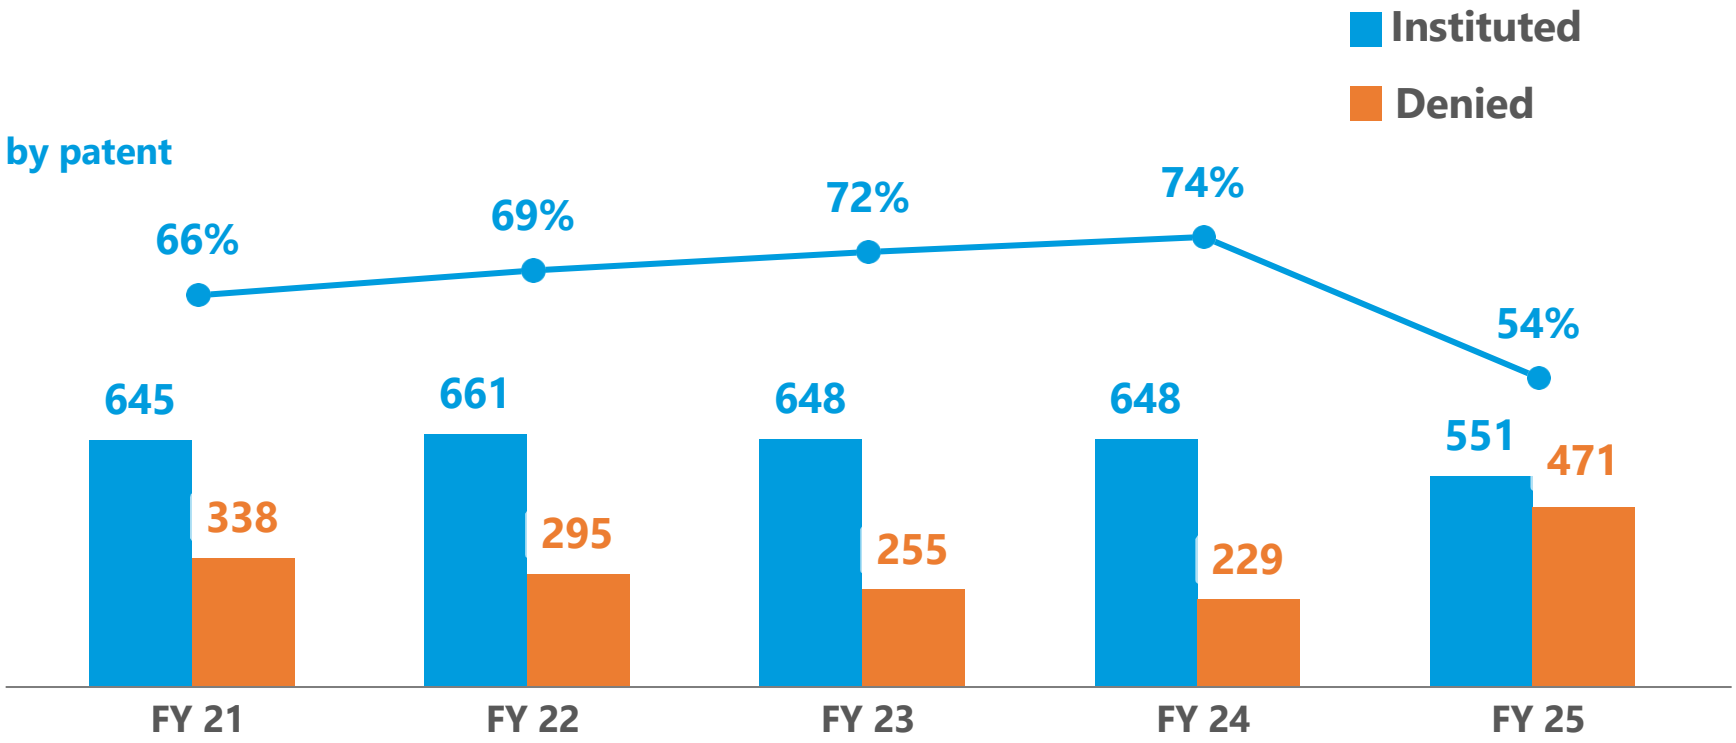
Institution rates by petition

FY 21 to FY 25: Oct. 1, 2020 to Sep. 30, 2025



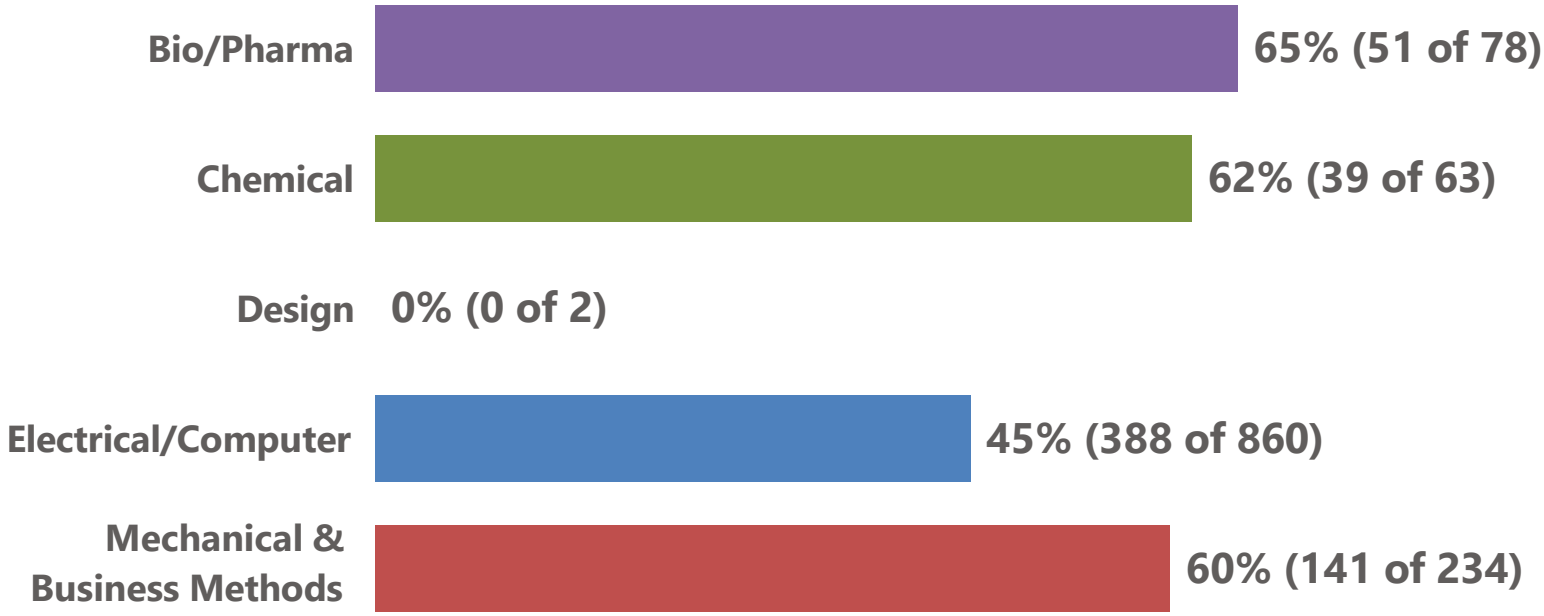
Institution rates by patent

FY 21 to FY 25: Oct. 1, 2020 to Sep. 30, 2025



Institution rates by technology

FY 25: Oct. 1, 2024 to Sep. 30, 2025

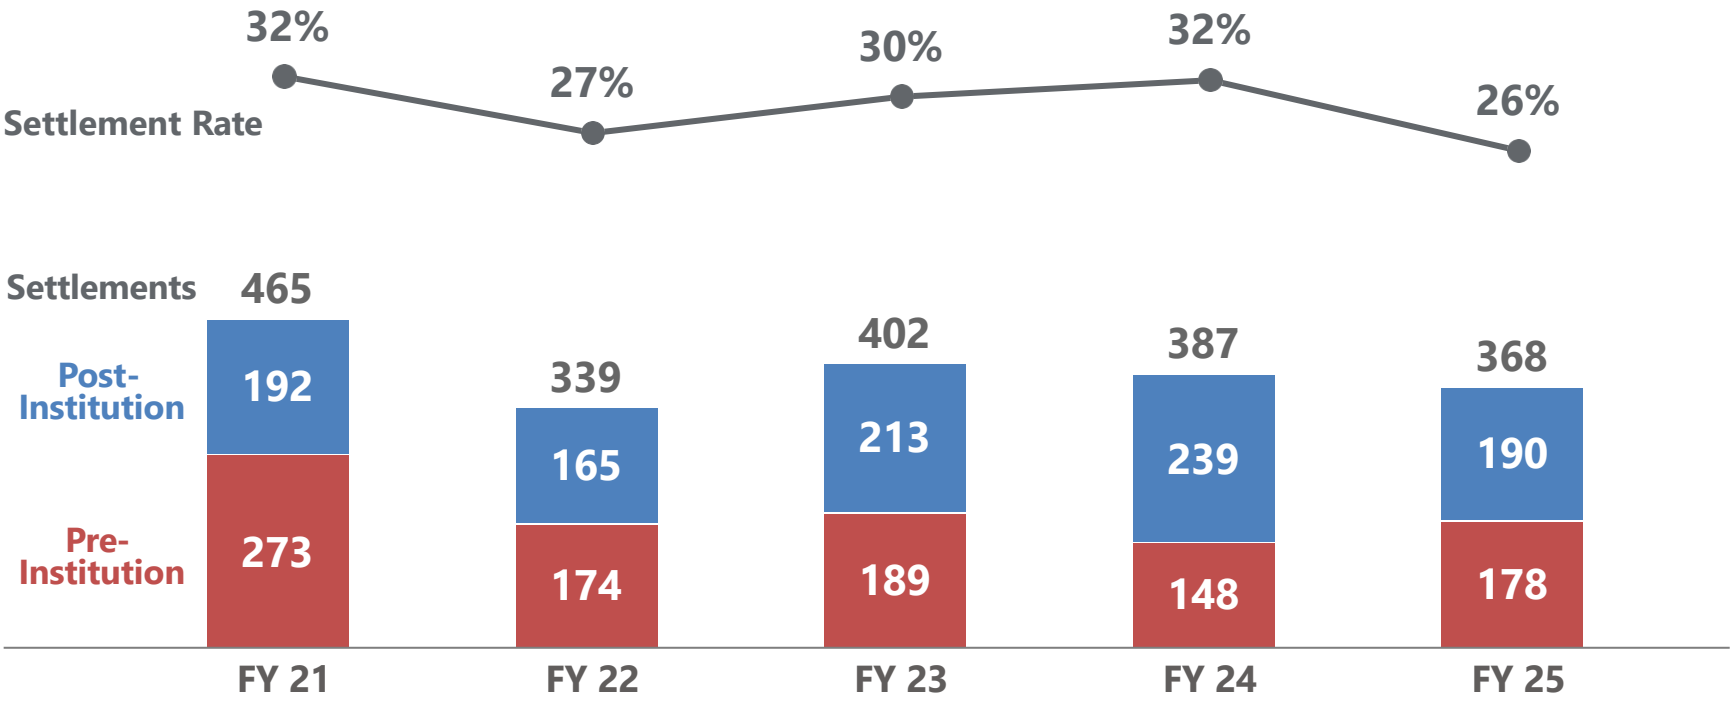


Institution rate for each technology is calculated by dividing petitions instituted by decisions on institution (i.e., petitions instituted plus petitions denied). The outcomes of decisions on institution responsive to requests for rehearing are excluded.

Motive Technologies Exhibit 2003, IPR2026-00034
Page 234

Settlements

FY 21 to FY 25: Oct. 1, 2020 to Sep. 30, 2025

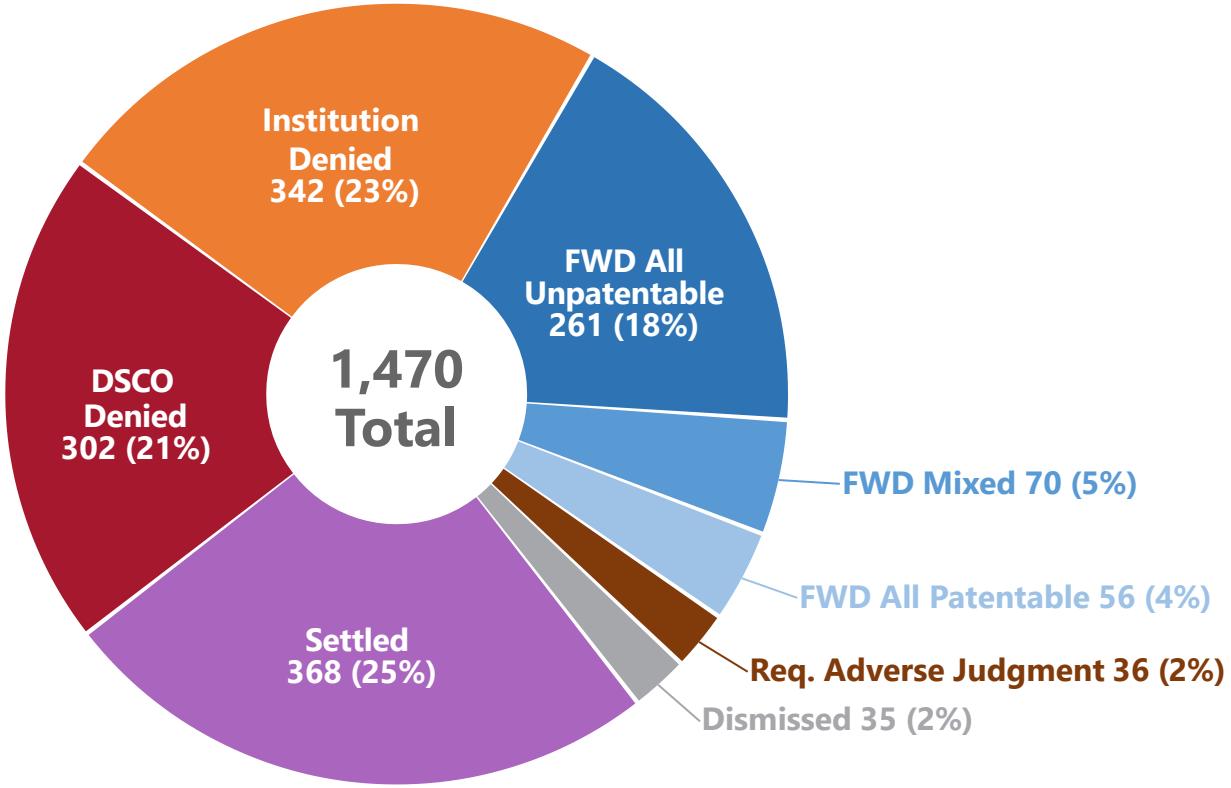


Settlement rate is calculated by dividing total settlements by concluded proceedings in each fiscal year (i.e., denied institution, settled, dismissed, requested adverse judgment, and final written decision), excluding joined cases.

Motive Technologies Exhibit 2003, IPR2026-00034
Page 235

Outcomes by petition

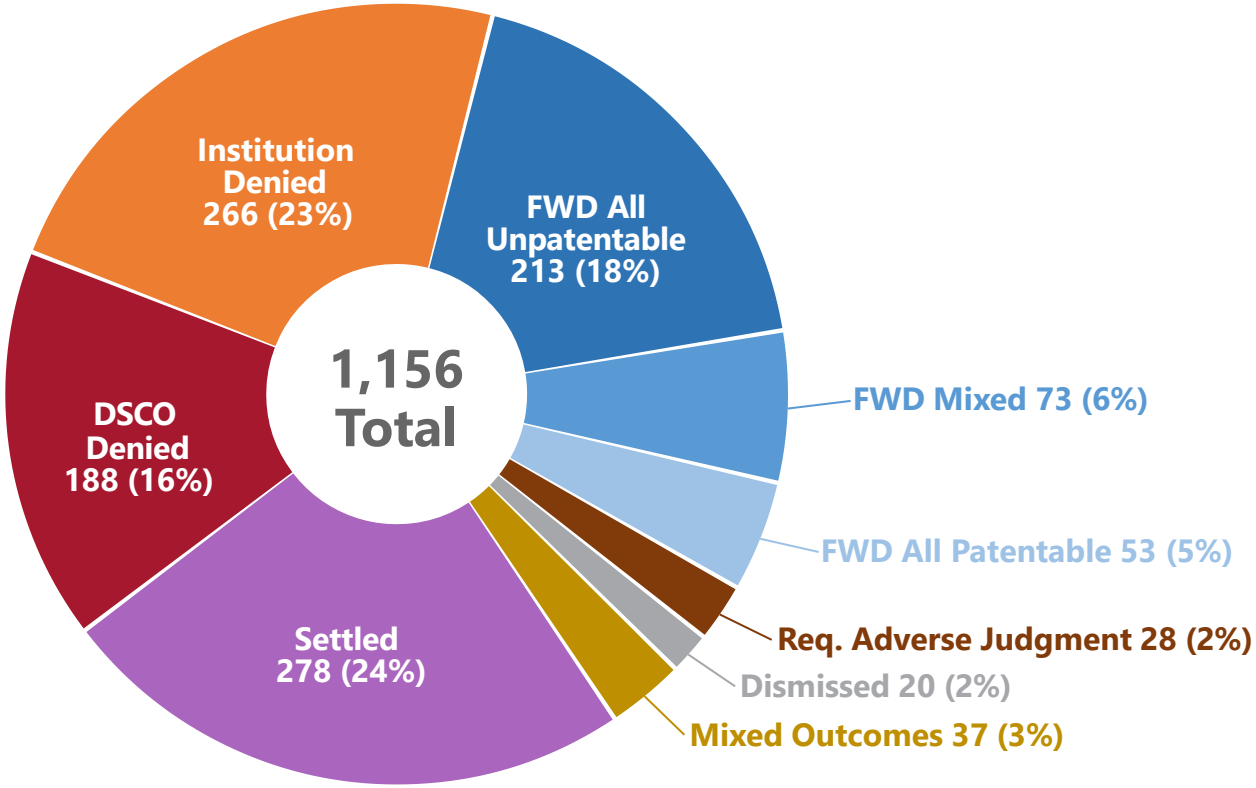
FY 25: Oct. 1, 2024 to Sep. 30, 2025



FWD patentability or unpatentability reported with respect to the claims at issue in the FWD. Joined cases are excluded.

Outcomes by patent

FY 25: Oct. 1, 2024 to Sep. 30, 2025

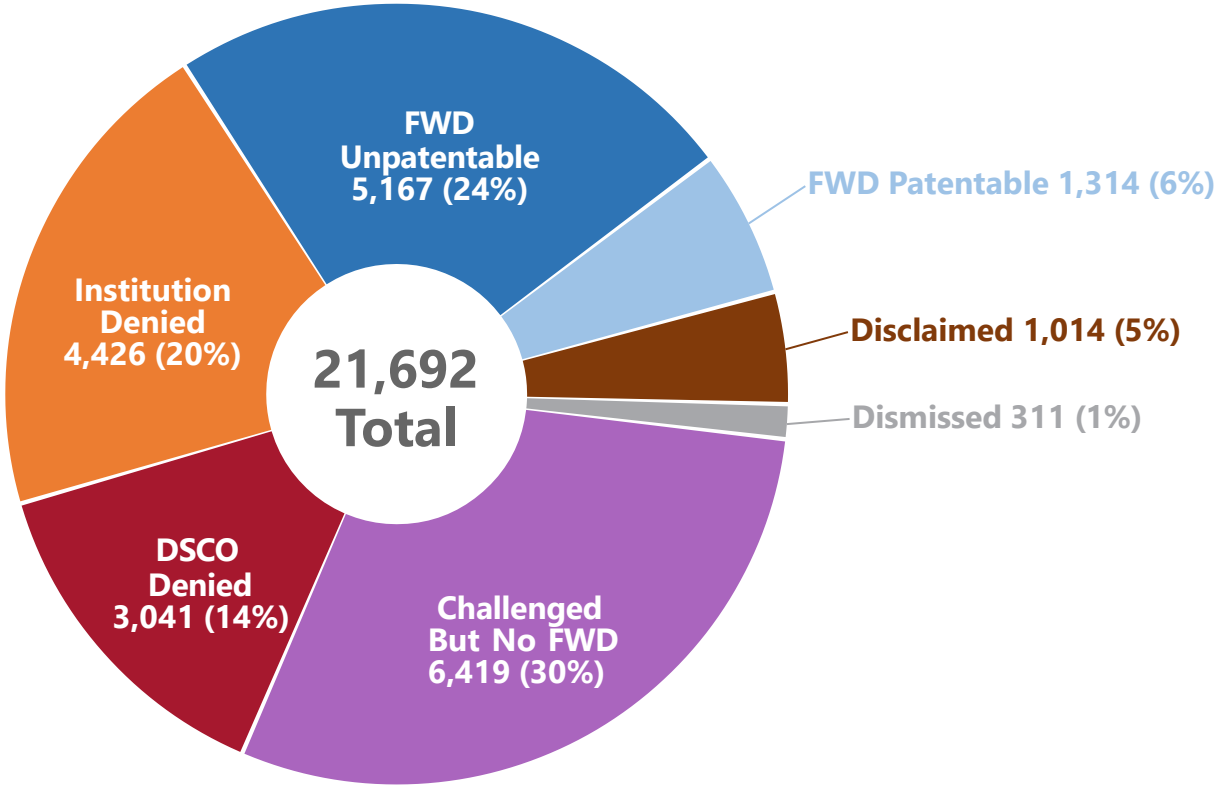


FWD patentability or unpatentability reported with respect to the claims at issue in the FWD. "Mixed Outcome" is shown for patents receiving more than one type of outcome from the list of: denied, settled, dismissed, and/or req. adverse judgement only. A patent is listed in a FWD category if it ever received a FWD, regardless of other outcomes.

Motive Technologies Exhibit 2003, IPR2026-00034
Page 237

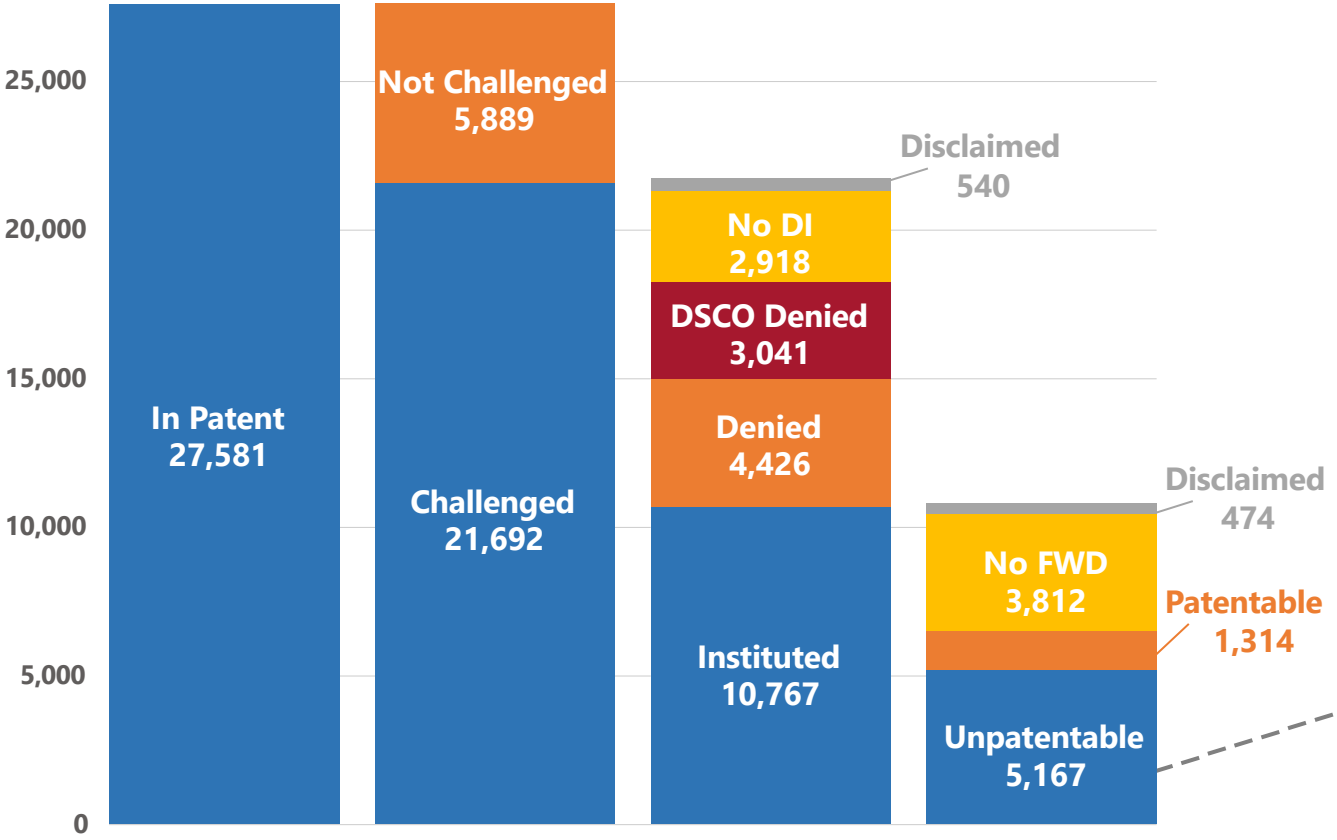
Outcomes by claim challenged

FY 25: Oct. 1, 2024 to Sep. 30, 2025



Outcomes by claim

FY 25: Oct. 1, 2024 to Sep. 30, 2025



"No DI" and "No FWD" mean the claim was challenged but not addressed in a DI/FWD, e.g., due to settlement.

24% of challenged claims and 48% of instituted claims were found unpatentable by a preponderance of the evidence.


UNITED STATES PATENT AND TRADEMARK OFFICE
DEPARTMENT OF COMMERCE

EXHIBIT D

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
18/527,226	12/01/2023	Ali HASSAN	189908.000302	7088

32361 7590 06/12/2024
GREENBERG TRAURIG (NY)
One Vanderbilt Avenue
NEW YORK, NY 10017

EXAMINER	
BROWN JR, HOWARD D	

ART UNIT	PAPER NUMBER
2488	

NOTIFICATION DATE	DELIVERY MODE
06/12/2024	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

cordesp@gtlaw.com
mendozae@gtlaw.com
njdocket@gtlaw.com

Office Action Summary	Application No. 18/527,226	Applicant(s) HASSAN et al.	
	Examiner HOWARD D BROWN JR	Art Unit 2488	AIA (FITF) Status Yes

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☑ Responsive to communication(s) filed on 12/01/2023.

☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2a)☐ This action is **FINAL.** 2b) ☑ This action is non-final.

3)☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

4)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

5) ☑ Claim(s) 21-40 is/are pending in the application.

5a) Of the above claim(s) _____ is/are withdrawn from consideration.

6) ☐ Claim(s) _____ is/are allowed.

7) ☑ Claim(s) 21-40 is/are rejected.

8) ☐ Claim(s) _____ is/are objected to.

9) ☐ Claim(s) _____ are subject to restriction and/or election requirement

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to **PPHfeedback@uspto.gov.**

Application Papers

10)☐ The specification is objected to by the Examiner.

11)☑ The drawing(s) filed on 12/01/2023 is/are: a)☑ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a)☐ All b)☐ Some** c)☐ None of the:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☑ Notice of References Cited (PTO-892)

2) ☑ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b) Paper No(s)/Mail Date 12/01/2023.

3) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date _____.

4) ☐ Other: _____.

DETAILED ACTION

1. This Office Action is sent in response to Applicant's communication received on 12/01/2023 for application number 18/527,226. The Office herby acknowledges receipt of the following and placed of record in file: Specification, Drawings, Abstract, Oath/Declaration, and claims.

Notice of Pre-AIA or AIA Status

2. The present application, filed on or after March 16, 2013, is being examined under the first inventor to file provisions of the AIA.

Information Disclosure Statement

3. The information disclosure statement (IDS) submitted on 12/01/2023 is in accordance with provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Double Patenting

4. The nonstatutory double patenting rejection is based on a judicially created doctrine grounded in public policy (a policy reflected in the statute) so as to prevent the unjustified or improper timewise extension of the "right to exclude" granted by a patent and to prevent possible harassment by multiple assignees. A nonstatutory double patenting rejection is appropriate where the conflicting claims are not identical, but at least one examined application claim is not patentably distinct from the reference claim(s) because the examined application claim is either anticipated by, or would have been obvious over, the reference claim(s). See, e.g., *In re Berg*, 140 F.3d 1428, 46 USPQ2d 1226 (Fed. Cir. 1998); *In re Goodman*, 11 F.3d 1046, 29 USPQ2d 2010 (Fed. Cir. 1993); *In re Longi*, 759 F.2d 887, 225 USPQ 645 (Fed. Cir. 1985); *In re Van*

Ornum, 686 F.2d 937, 214 USPQ 761 (CCPA 1982); *In re Vogel*, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); *In re Thorington*, 418 F.2d 528, 163 USPQ 644 (CCPA 1969).

5. A timely filed terminal disclaimer in compliance with 37 CFR 1.321(c) or 1.321(d) may be used to overcome an actual or provisional rejection based on nonstatutory double patenting provided the reference application or patent either is shown to be commonly owned with the examined application, or claims an invention made as a result of activities undertaken within the scope of a joint research agreement. See MPEP § 717.02 for applications subject to examination under the first inventor to file provisions of the AIA as explained in MPEP § 2159. See MPEP §§ 706.02(l)(1) - 706.02(l)(3) for applications not subject to examination under the first inventor to file provisions of the AIA. A terminal disclaimer must be signed in compliance with 37 CFR 1.321(b).

6. The USPTO Internet website contains terminal disclaimer forms which may be used. Please visit www.uspto.gov/patent/patents-forms. The filing date of the application in which the form is filed determines what form (e.g., PTO/SB/25, PTO/SB/26, PTO/AIA/25, or PTO/AIA/26) should be used. A web-based eTerminal Disclaimer may be filled out completely online using web-screens. An eTerminal Disclaimer that meets all requirements is auto-processed and approved immediately upon submission. For more information about eTerminal Disclaimers, refer to www.uspto.gov/patents/process/file/efs/guidance/eTD-info-I.jsp.

7. Claims 21-22, 28 are rejected on the ground of nonstatutory double patenting as being unpatentable over claims 1, 5, 8 of U.S. Patent No. 11,875,580. Although the claims at issue are not identical, they are not patentably distinct from each other: Table 1 shows comparison between the instant claims and the U.S. Patent 11,875,580 B1 claims.

See Table:

App: 18/527,226	Patent: 11,875,580
21. (New) A method comprising: receiving an image of a roadway recorded by a camera device installed within a vehicle; detecting a horizon line in the image; overlaying a line on the image to generate an overlaid image; transmitting the overlaid image to a computing device over a network; receiving a modification of the line from the computing device, the modification comprising a new line at a second position; computing a camera parameter based on the new line; and transmitting data representing the camera parameter to the camera device.	1. A method comprising: receiving, over a network from a camera device, a video comprising a set of image frames; identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line; computing at least one camera parameter based on the one or more lines; overlaying the one or more lines on the video to generate an overlaid video; transmitting the overlaid video to an annotator device for manual review; receiving a confirmation from the annotator device, the confirmation indicating that the one or more lines are accurate; and transmitting data representing the at least one camera parameter to the camera device.
22. (New) The method of claim 21, wherein the camera parameter comprises one of camera height, viewing angle, and road plane normal.	5. The method of claim 1, wherein the at least one camera parameter comprises a camera height and road plane normal.

28. (New) A non-transitory computer-readable storage medium for tangibly storing computer program that, when executed by a computer processor, perform operations of: receiving an image of a roadway recorded by a camera device installed within a vehicle; detecting a horizon line in the image; overlaying a line on the image to generate an overlaid image; transmitting the overlaid image to a computing device over a network; receiving a modification of the line from the computing device, the modification comprising a new line at a second position; computing a camera parameter based on the new line; and transmitting data representing the camera parameter to the camera device.	8. A method comprising: receiving a video, the video comprising one or more image frames; generating a set of intermediate features representing the video using a backbone network; inputting the set of intermediate features into a lane or horizon prediction head, the lane or horizon prediction head comprising a neural network; predicting a horizon line using the lane or horizon prediction head, the horizon line comprising a set of horizon markers; and using the horizon line to compute at least one camera parameter.

Claim Rejections - 35 USC § 103

8. The following is a quotation of 35 U.S.C. 103 which forms the basis for all obviousness rejections set forth in this Office action:

> A patent for a claimed invention may not be obtained, notwithstanding that the claimed invention is not identically disclosed as set forth in section 102, if the differences between the claimed invention and the prior art are such that the claimed invention as a whole would have been obvious before the effective filing date of the claimed invention to a person having ordinary skill in the art to which the claimed invention pertains. Patentability shall not be negated by the manner in which the invention was made.

9. The factual inquiries for establishing a background for determining obviousness under 35 U.S.C. 103 are summarized as follows:

1. Determining the scope and contents of the prior art.

2. Ascertaining the differences between the prior art and the claims at issue.

3. Resolving the level of ordinary skill in the pertinent art.

4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

10. This application currently names joint inventors. In considering patentability of the claims the examiner presumes that the subject matter of the various claims was commonly owned as of the effective filing date of the claimed invention(s) absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and effective filing dates of each claim that was not commonly owned as of the effective filing date of the later invention in order for the examiner to consider the applicability of 35 U.S.C. 102(b)(2)(C) for any potential 35 U.S.C. 102(a)(2) prior art against the later invention.

11. Claim(s) 21-40 is/are rejected under 35 U.S.C. 103 as being unpatentable over Kwant1 et al., [US Pub. No.: 2019/0034740 A1] Kwant2 et al., [US Pub. No.: 2019/0102674 A1].

Re. Claim 21. Kwant1 discloses

A method comprising: receiving an image of a roadway recorded by a camera device installed within a vehicle ***[network for receiving image signals from a vehicle camera of a roadway |Fig.1 el 113, 0076, 0081]***;

detecting a horizon line in the image ***[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 |Abstract]***;

overlaying a line on the image to generate an overlaid image ***[Fig.4 401A-C]***;

transmitting the overlaid image to a computing device over a network ***[In step 607, the horizon estimation platform 105 labels or annotates one or more map features depicted in the image based on the horizon|Fig.1 el 100, Fig.6 Step 607, 0071]***;

In the same field of endeavor Kwant2 discloses:

receiving a modification of the line from the computing device, the modification comprising a new line at a second position ***[The learner module then computes an accuracy of the predictions for the initial set of model parameters. If the accuracy or level of performance does not meet a threshold or configured level, the learner module incrementally adjusts the model parameters until the model generates predictions at a desired or configured level of accuracy with respect to the manually annotated labels in the training data (e.g., the ground truth data).|0051]***;

computing a camera parameter based on the new line ***[The machine learning system 103 generates respective distributions or probability models for the training set 401 and the reservoir set 405 according to the various embodiments described herein. These distributions are combined into a single model (e.g., the sampling distribution 401) describing a candidate observation 407's similarity to the training set and/or the candidate pool at evaluation time to either accept or reject the candidate observation 407 from the reservoir set 405 for annotation and inclusion in the training data set 403]***;

and transmitting data representing the camera parameter to the camera device ***[data representing camera parameter to the camera device |Fig. 8a]***.

Therefore, it would have been obvious to one with ordinary skill in the art before the effective filing date of the claimed invention to combine the horizon estimation using lane lines of Kwant1 with Kwant2 for selecting representative training observations for annotation in machine learning model.

Re. Claim 22, Kwant1 discloses:

The method of claim 21, wherein the camera parameter comprises one of camera height, viewing angle, and road plane normal ***[camera height, viewing angle and roadway parameters |Fig.3, 0051,0082]***.

Re. Claim 23, Kwant1 discloses:

The method of claim 21, wherein receiving the image of the roadway comprises receiving an image of the roadway while the vehicle is moving ***[images roadway are received during vehicle operation |Fig.4a-b, 5a-c]***.

Re. Claim 24, Kwant1 discloses:

The method of claim 23, wherein receiving the image of the roadway while the vehicle is moving comprises detecting that the vehicle is traveling above a pre-defined speed or for a pre-defined duration ***[vehicle speed and roadway speed level information is tracked and or received |0099, 0101]***.

Re. Claim 25, Kwant1 discloses:

The method of claim 21, wherein transmitting the overlaid image to the computing device over the network comprises transmitting the overlaid image to a web-based application ***[components of the horizon estimation system capable to transmit images over a network to a web-based application |Fig.1, 0038, 0056, 0063, 0081]***.

Re. Claim 26, Kwant1 discloses:

The method of claim 21, wherein receiving the modification of the line from the computing device comprises receiving an identification of a new horizon line by a user of the computing device ***[to account for variances in the slopes of the roadway, the horizon estimation platform 105 detects a slope of the roadway, and adjusts the horizon based on the slope. The horizon estimation platform 105, for instance, can determine the slope from one or more sensors of a vehicle… to the examiner adjusting horizon line is identifying a new horizon based captured sensor data |0057]***.

Re. Claim 27, the rejection of claim 1 is incorporated herein.

Kwant1 & Kwant2 discloses the claim limitations, as follows:

The method of claim 21, wherein transmitting the data representing the camera parameter to the camera device ***[Kwant2: data representing camera parameter to the camera device |Fig. 8a]*** further comprises recomputing the camera parameter based on the modification of the line ***[Kwant1: a parametric representation of a line can have 12 parameters in one set to describe a single lane line. |0013]***.

Re. Claim 28, This claim is interpreted and rejected for the same reason set forth in claim 21.

Re. Claim 29, This claim is interpreted and rejected for the same reason set forth in claim 22.

Re. Claim 30, This claim is interpreted and rejected for the same reason set forth in claim 23.

Re. Claim 31, This claim is interpreted and rejected for the same reason set forth in claim 24.

Re. Claim 32, This claim is interpreted and rejected for the same reason set forth in claim 25.

Re. Claim 33, This claim is interpreted and rejected for the same reason set forth in claim 26.

Re. Claim 34, This claim is interpreted and rejected for the same reason set forth in claim 27.

Re. Claim 35, This claim is interpreted and rejected for the same reason set forth in claim 1 & 28 including a device comprising: a processor ***[Fig. 14, box 1403]***;
and a storage medium for tangibly storing thereon program logic for execution by the processor ***[Fig. 14 box 1405].***

Re. Claim 36, This claim is interpreted and rejected for the same reason set forth in claims 23 & 30.

Re. Claim 37, This claim is interpreted and rejected for the same reason set forth in claims 24 & 31.

Re. Claim 38, This claim is interpreted and rejected for the same reason set forth in claims 25 & 32.

Re. Claim 39, This claim is interpreted and rejected for the same reason set forth in claims 26 & 33.

Re. Claim 40, This claim is interpreted and rejected for the same reason set forth in claims 27 & 34.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to HOWARD D BROWN JR whose telephone number is (571)272-4371. The examiner can normally be reached Monday - Friday 7:30AM - 5:00PM EST.

Examiner interviews are available via telephone, in-person, and video conferencing using a USPTO supplied web-based collaboration tool. To schedule an interview, applicant is encouraged to use the USPTO Automated Interview Request (AIR) at http://www.uspto.gov/interviewpractice.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Sathyanarayanan Perungavoor can be reached on 5712727455. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of published or unpublished applications may be obtained from Patent Center. Unpublished application information in Patent Center is available to registered users. To file and manage patent submissions in Patent Center, visit:

https://patentcenter.uspto.gov. Visit https://www.uspto.gov/patents/apply/patent-center for more information about Patent Center and https://www.uspto.gov/patents/docx for information about filing in DOCX format. For additional questions, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

HOWARD D. BROWN JR
Primary Examiner
Art Unit 2488

/HOWARD D BROWN JR/
Examiner, Art Unit 2488

EXHIBIT E

MEMORANDUM

To: All PTAB Judges

From: John A. Squires
Under Secretary of Commerce for Intellectual Property and
Director of the United States Patent and Trademark Office

Subject: Director Institution of AIA Trial Proceedings

Date: October 17, 2025

To improve efficiency, consistency, and adherence to the statutory requirements for institution of trial, effective October 20, 2025, the Director will determine whether to institute trial for *inter partes* review ("IPR") and post-grant review ("PGR") proceedings.[1] This process will maintain PTAB's capacity to conduct IPR and PGR trials and promote consistent application of considerations for institution of trial proceedings before the PTAB. This approach to institution flows from the processes outlined in the March 26, 2025 memorandum entitled "Interim Processes for PTAB Workload Management" ("Interim Processes"),[2] under which the Director determines whether or not to deny a petition based on discretionary considerations.

Similar to the discretionary considerations process, the Director, in consultation with at least three PTAB judges, will determine whether to institute trials in all IPR and PGR proceedings. Upon review of discretionary considerations, the merits, and non-discretionary

---

[1] Congress provided that the Director determines whether to institute trials under the America Invents Act. *See* 35 U.S.C. § 314(a) ("The Director may not authorize an inter partes review to be instituted unless the Director determines that the information presented in the petition filed under section 311 and any response filed under section 313 shows that there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition."); *id.* § 314(b) ("The Director shall determine whether to institute an inter partes review under this chapter pursuant to a petition . . . ."); *id.* § 314(c) ("The Director shall notify the petitioner and patent owner, in writing, of the Director's determination under subsection (a), and shall make such notice available to the public as soon as is practicable."); *see also id.* § 324(a), (c), (d) (similar).

[2] Available at https://www.uspto.gov/sites/default/files/documents/InterimProcesses-PTABWorkloadMgmt-20250326.pdf.

considerations, if the Director determines that institution is appropriate on at least one ground for one challenged claim, the Director will issue a summary notice to the parties granting institution. *See* 35 U.S.C. §§ 314(c), 324(d). Similarly, if the Director determines that institution is not appropriate, whether based on discretionary considerations, the merits, or other non-discretionary considerations, the Director will issue a summary notice denying institution. In proceedings involving novel or important factual or legal issues, the Director may issue a decision on institution addressing those issues. Additionally, where the Director determines detailed treatment of issues raised in a petition is appropriate (e.g., complex claim construction issues, priority analysis, or real party in interest determination), the Director may refer the decision on institution to one or more members of the PTAB. The Office has issued more than 580 decisions under the Interim Processes, providing substantial guidance on how the Director will handle discretionary considerations. Any instituted IPR or PGR proceeding will be referred to a three-member panel of the PTAB to conduct the trial and that panel will be assigned according to PTAB Standard Operating Procedure (SOP) 1 (Rev. 16).[3]

This Memorandum supersedes the Interim Processes to the extent that (1) routine decisions on institution will be limited to summary notices, and (2) merit-based decisions on whether to institute petitions will not be referred to a three-member panel of the PTAB. The process for briefing discretionary considerations, as outlined in the Interim Processes and the Discretionary Decisions webpage,[4] and the process for briefing the merits and non-statutory considerations will remain the same. Further, all petitions referred to the PTAB for consideration of the merits and non-discretionary considerations under the Interim Processes prior to October 20, 2025 will remain with a three-member panel.

---

[3] Available at https://www.uspto.gov/sites/default/files/documents/sop1_r16_final.pdf.
[4] Available at https://www.uspto.gov/patents/ptab/interim-director-discretionary-process.

Motive Technologies Exhibit 2003, IPR2026-00034
Page 257

GIBSON, DUNN & CRUTCHER LLP
JOSH KREVITT (SBN 208552)
JKrevitt@gibsondunn.com
STUART M. ROSENBERG (SBN 239926)
SRosenberg@gibsondunn.com
310 University Avenue
Palo Alto, CA 94301
Telephone: 650.849.5300
Facsimile: 650.849.5333

BRIAN A. ROSENTHAL (*pro hac vice*)
BRosenthal@gibsondunn.com
KATHERINE Q. DOMINGUEZ (*pro hac vice*)
KDominguez@gibsondunn.com
ALLEN KATHIR (*pro hac vice*)
AKathir@gibsondunn.com
ALLYSON PARKS (*pro hac vice*)
AParks@gibsondunn.com
200 Park Avenue
New York, NY 10166
Telephone: 212.351.4000
Facsimile: 212.351.4035

GIBSON, DUNN & CRUTCHER LLP
JAYSEN S. CHUNG (SBN 280708)
JSChung@gibsondunn.com
One Embarcadero Center, Suite 2600
San Francisco, CA 94111
Telephone: 415.393.8200
Facsimile: 415.393.8306

NATHANIEL R. SCHARN (SBN 305836)
NScharn@gibsondunn.com
3161 Michelson Drive, Suite 1200
Irvine, CA 92612
Telephone: 949.451.3800
Facsimile: 949.451.4220

HANNAH L. BEDARD (*pro hac vice*)
HBedard@gibsondunn.com
WENDY W. CAI (*pro hac vice*)
WCai@gibsondunn.com
1700 M STREET, NW
Washington, DC 20036
Telephone: 202.955.8500
Facsimile: 202.467.0539

Attorneys for Defendant Samsara Inc.

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

MOTIVE TECHNOLOGIES, INC.,

Plaintiff and Counterclaim Defendant,

v.

SAMSARA INC.,

Defendant and Counterclaim Plaintiff.

CASE NO. 3:24-CV-00902-JD

[PROPOSED] ORDER GRANTING SAMSARA, INC.'S MOTION TO STAY MOTIVE'S PATENT INFRINGEMENT CLAIMS PENDING *INTER PARTES* REVIEW

Defendant and Counterclaim Plaintiff Samsara, Inc. ("Samsara") moves for entry of an order partially staying this case in relation to the First, Ninth, and Tenth Causes of Action of Plaintiff and Counterclaim Defendant Motive Technologies Inc. ("Motive") (Dkt. 107, ¶¶157-69, 249-74, "Motive Patent Claims") pending final resolution of the *Inter Partes* Review proceedings concerning U.S. Patent Nos. 11,875,580; 12,062,243; and 12,136,276 ("IPR Proceedings"), including any appeals arising from those proceedings. The Court has considered the submissions and arguments of the parties.

IT IS on this ___ day of _____, 2025,

ORDERED that Samsara's Motion to Stay the Motive Patent Claims pending the IPR Proceedings is hereby GRANTED; and it is further

ORDERED that the part of this case in relation to the Motive Patent Claims is stayed for the duration of IPR Proceedings and any appeals from those proceedings; and it is further

ORDERED that all deadlines in the Scheduling Order (Dkt. No. 88) regarding the Motive Patent Claims are hereby vacated.

_____

Hon. James Donato, U.S.D.J.