

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSARA INC.
Petitioner

v.

MOTIVE TECHNOLOGIES, INC.
Patent Owner

Case IPR2025-00574
U.S. Patent No. 11,875,580

DECLARATION OF DR. TREVOR J. DARRELL

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

Samsara EX1085
Samsara v. Motive
IPR2025-00574

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	PERSON OF ORDINARY SKILL IN THE ART	4
	A. The automotive industry has a long history with computer vision.....	6
	B. Motive’s POSA definition adds additional skill to the POSA.....	9
III.	I HAVE EXTENSIVE EXPERIENCE WITH COMPUTER VISION FOR THE AUTOMOTIVE INDUSTRY.....	10
IV.	I DISAGREE WITH MOTIVE’S IMPLICIT CONSTRUCTION OF “PREDICTIVE MODEL”	14
	A. The ’580 patent does not exclude a geometric approach.....	14
	B. The prosecution history does not support a narrow construction of “predictive model.”	15
V.	GROUND 1: CHOE RENDERS OBVIOUS CLAIMS 1-4.	22
	A. Choe discloses the “predictive model.”	22
	B. The calibration factors disclosed in Choe are or are used to compute “camera parameters.”	25
VI.	GROUND 2: CHOE IN VIEW OF KHAN/ALI RENDERS OBVIOUS CLAIMS 1-5.....	26
	A. My opinion on the motivation to combine Choe and Khan/Ali has not changed.....	26
	B. Choe in view of Khan/Ali discloses the “predictive model.”	30
	1. Khan and Ali’s CNN is a predictive model that is used to identify lane lines and the horizon line.	31
	2. Khan/Ali’s horizon line estimation is a predictive model.	38
	3. Khan and Ali’s end2end algorithm for detecting lane lines and the horizon line is a predictive model.....	40
	4. Khan/Ali relaxes the requirement of having three lane lines in frame.	44
	C. Choe and Khan/Ali perform calibration based on video.....	46
VII.	GROUND 3: MY OPINION ON THE MOTIVATION TO COMBINE CHOE AND WORKMAN HAS NOT CHANGED.....	48

VIII. GROUND 4: MY OPINION ON THE MOTIVATION TO COMBINE CHOE IN VIEW OF WESTMACOT HAS NOT CHANGED.	53
IX. GROUNDS 5 AND 6: MY OPINION ON THE MOTIVATION TO COMBINE HAS NOT CHANGED	59
X. CONCLUSION.....	60

I, Trevor J. Darrell, hereby declare as follows.

I. INTRODUCTION

1. I have been retained on behalf of Samsara Inc. for the above-captioned *inter partes* review proceeding. I understand that this proceeding involves U.S. Patent No. 11,875,580 (EX1001, “the ’580 patent”), which is titled “Camera Initialization for Lane Detection and Distance Estimation Using Single-View Geometry.”

2. I previously submitted a Declaration (EX1003) in support of the Petition in this proceeding. My qualifications are provided in my prior declaration, and my CV was submitted as Exhibit 1004.

3. Since my prior declaration, I have reviewed and considered the following additional materials:

Paper	Description
12	Decision Granting Institution (“DI”)
20	Patent Owner’s Response (“POR”)

Exhibit	Description
2003	Declaration of Scott Andrews (June 18, 2025)
2003 ¹	Declaration of Scott Andrews (Nov. 17, 2025)
1061	Deposition Transcript of Scott Andrews

¹ Mr. Andrews submitted two declarations with the same exhibit number, EX2003. To distinguish between the two, I will refer to his first (POPR) declaration as “POPR EX2003” and the second, “revised” (POR) declaration as “EX2003.”

1062	IEEE Intelligent Vehicles Symposium, <i>available at</i> https://ieeetitss.org/conf/iv/
1063	Mobileye About, <i>available at</i> https://www.mobileye.com/about/ [Including PDF printed pages of webpages and screen shots for clarity]
1064	<i>Intel buys driverless car technology firm Mobileye</i> , BBC News (Mar. 13, 2017), <i>available at</i> https://www.bbc.co.uk/news/business-39253422
1065	<i>BMW Group, Intel and Mobileye Team Up to Bring Fully Autonomous Driving to Streets by 2021</i> , BMW Group (Jan. 7, 2016), <i>available at</i> https://www.press.bmwgroup.com/global/article/detail/T0261586EN/bmw-group-intel-and-mobileye-team-up-to-bring-fully-autonomous-driving-to-streets-by-2021
1066	<i>The Road to Autonomous Driving: How BMW and Mobileye are Crowdsourcing Real-Time Data</i> , fxis.ai (Sept. 6, 2014), <i>available at</i> https://fxis.ai/edu/the-road-to-autonomous-driving-how-bmw-and-mobileye-are-crowdsourcing-real-time-data/
1067	F. Wu et al., “Decentralized Vehicle Coordination: The Berkeley DeepDrive Drone Dataset and Consensus-Based Models,” 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 2025, pp. 2438-2444, doi: 10.1109/ICRA55743.2025.11127341. (https://ieeexplore.ieee.org/abstract/document/11127341)
1068	Fisher Yu and Haofeng Chen and Xin Wang and Wenqi Xian and Yingying Chen and Fangchen Liu and Vashisht Madhavan and Trevor Darrell, “BDD100K: Images 100K Dataset,” (2020) <i>available at</i> https://datasetninja.com/bdd100k
1069	F. Yu et al., “BDD100K: A Diverse Driving Dataset for Heterogeneous Multitask Learning,” 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 2020, pp. 2633-2642, doi: 10.1109/CVPR42600.2020.00271.
1070	Vijay Govindarajan, “Adaptive Prediction and Planning for Safe and Effective Autonomous Vehicles,” EECS Department, University of California, Berkeley, Technical Report No. UCB/EECS-2021-263, December 17, 2021, https://www2.eecs.berkeley.edu/Pubs/TechRpts/2021/EECS-2021-263.html
1071	Yang Gao and Trevor Darrell, “End to End Learning in Autonomous Driving Systems,” EECS Department, University of California,

	Berkeley, Technical Report No. UCB/EECS-2020-5, January 8, 2020, https://www2.eecs.berkeley.edu/Pubs/TechRpts/2020/EECS-2020-5.html
1072	NVIDIA Video (December 7, 2016), <i>available at</i> https://www.youtube.com/watch?v=eeR5xDaNCJ8&pp=ygUOdHJldm9yIGRhcnJlbGw%3D
1073	NVIDIA Webpage (December 7, 2016), <i>available at</i> https://www.youtube.com/watch?v=eeR5xDaNCJ8&pp=ygUOdHJldm9yIGRhcnJlbGw%3D
1074	2017 BEARS: Panel 1 Video (February 21, 2017), <i>available at</i> https://www.youtube.com/watch?v=ltbufoOON8Y
1075	2017 BEARS: Panel 1 Webpage (February 21, 2017), <i>available at</i> https://www.youtube.com/watch?v=ltbufoOON8Y
1076	Harris, Mark, <i>The AI Dashcam App That Wants to Rate Every Driver in the World</i> , IEEE Spectrum (Jun. 15, 2016), <i>available at</i> https://spectrum.ieee.org/the-ai-dashcam-app-that-wants-to-rate-every-driver-in-the-world
1077	Trevor Darrell – EECS at UC Berkeley, <i>available at</i> https://www2.eecs.berkeley.edu/Faculty/Homepages/darrell.html
1078	PATH Directory, https://path.berkeley.edu/people/path-directory/
1079	About Us – PATH, https://path.berkeley.edu/about-us/
1080	PATH – Institute of Transportation Studies, <i>available at</i> https://its.berkeley.edu/research-centers/path
1081	Berkeley DeepDrive, <i>available at</i> https://deepdrive.berkeley.edu/
1082	Berkeley DeepDrive – Institute of Transportation Studies, <i>available at</i> https://its.berkeley.edu/research-centers/berkeley-deepdrive
1083	Mobileye Algorithms Developers, <i>available at</i> https://web.archive.org/web/20161228222811/http://careers.mobileye.com/Job-Search/R-D-Job-Description/Algorithms-Developers/28375.aspx [Including PDF printed pages of webpages and screen shots for clarity]
1084	Ahmed Elgammal, <i>CS534 – Computer Vision Camera Calibration</i> , Rutgers University, <i>available at</i> https://people.cs.rutgers.edu/~elgammal/classes/cs534/lectures/Calibration.pdf

4. I also considered all other materials cited herein. My statements in my prior declaration regarding my review of the '580 patent and related materials

remain unchanged, as do my understanding of the relevant legal principles.

5. I understand that Patent Owner Motive Technologies, Inc. (“Motive” or “Patent Owner”) submitted a Patent Owner Response and a declaration from Mr. Scott Andrews to support the arguments in its Patent Owner Response. I have been asked to address certain arguments made by Motive, as well as related testimony from Mr. Andrews. My work on this case is being billed at my standard hourly rate of \$810 per hour. My compensation does not depend on the outcome of this *inter partes* review and in no way impacts the substance of my statements in this declaration.

II. PERSON OF ORDINARY SKILL IN THE ART

6. I understand that Motive and Mr. Andrews dispute the definition of the Person of Ordinary Skill in the Art (“POSA”). As the following chart shows, Motive argues that a POSA should include 1-2 years training and/or experience in design of software for camera/sensor systems in the automotive industry whereas the definition in my previous declaration required two years of experience and was not limited to the automotive industry.

Definition in My Previous Declaration	Definition in Motive’s POR
Someone with “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)	“A POSITA for the ’580 Patent as of 2021 would have a bachelor’s degree in engineering or computer science with one- or two-years training and/or experience in design of software for camera/sensor systems in the automotive industry.” POR, 8-9.

classical machine/computer vision algorithms (e.g., edge detection, line identification, computation of camera parameters from horizon line, etc.). Additional education could serve as a substitute for the experience requirement.” EX1003, ¶77.	
--	--

7. I disagree with Motive’s definition and maintain the definition in my previous declaration. Although the preferred embodiment relates to automotive systems, the ’580 patent claims are not so limited. The machine learning systems that predict lines in a video, as described in the ’580 patent, have broader applications than just the automotive industry. Indeed, the inventors of the ’580 patent presented their work at computer vision conferences that were not limited to the automotive industry. *See* EX1052-1056 (presenting results at IEEE Winter Conference on Applications of Computer Vision (WACV)).

8. However, if the Board adopts Motive’s definition of a POSA, that definition would not materially impact my opinions. As I explain in Section II.A, an engineer with one- or two-years training and/or experience in the design of software for camera/sensor systems in the automotive industry would have been familiar with computer vision, including the machine learning methods and machine/computer vision algorithms I specified in my POSA definition.

9. Further, as I explain in Section II.B, adding automotive industry experience to my POSA definition would *raise* the level of skill in the art because

it would add additional industry experience that my definition did not require. As I understand from counsel, it is generally easier to establish obviousness under a higher level of ordinary skill in the art.

A. The automotive industry has a long history with computer vision.

10. Before the October 4, 2021, priority date of the '580 patent, there was ample activity in the automotive industry regarding the use of computer vision/machine learning methods and algorithms for automotive sensor design. The automotive industry was aware that computer vision and machine learning were critical to the establishment of autonomous driving.

11. For example, IEEE's Intelligent Vehicles Symposium (IV) has been occurring since 1992. EX1062. In 2016, I went to IV in Gothenburg, Sweden to network with automotive companies that were attending the conference, including Nexar, an advanced driver assistance systems (ADAS) dashcam company. At the conference, paper presentations included (a) 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, computerization of camera parameters from horizon line, etc.). As this example shows, the automotive industry was well aware of these machine learning methods and classical machine/computer vision algorithms before 2021.

12. It was also quite common for automotive companies to attend

computer vision conferences, such as the ComputerVision Foundation's WACV.

For example, the inventors of the '580 patent presented the Khan and Ali references at the 2020 WACV. *See* EX1053; EX1054; EX1055; EX1056.

13. Another example of computer vision's importance to automotive sensors and cameras is Mobileye Global Inc. (Mobileye). In 1999, the founder of Mobileye, Amnon Shashua, took his academic research in computer vision and applied it to monocular vision systems for cars to detect other vehicles. EX1063. This was done through using a camera and processor. EX1063. In 2001, Mobileye designed a custom chip specifically for running the computer vision algorithm. EX1063. By around the mid-2010s, Mobileye was considered an industry leader in ADAS, in part for their use of a single camera and specialized processor. In 2017, Mobileye was acquired by Intel for \$15.3 billion. EX1064. By that point, Mobileye had contracts with at least 27 car manufacturers and was working with BMW to develop an autonomous driving system. EX1064; EX1065; EX1066. A POSA would have therefore known of Mobileye's influence on the industry and the industry-wide interest in the application of computer vision technology in automotive sensors and cameras.

14. It is also my opinion that someone with a bachelor's degree in engineering or computer science with one- or two-years training and/or experience in design of software for camera/sensor systems in the automotive industry would

have known about and implemented the computer vision methods and classical machine/computer vision algorithms I detailed in my POSA definition. For example, at Mobileye, an entry level position requiring a Bachelor's of Science in Computer Science or similar field (e.g., an Algorithm Developer) would have been prototyping solutions for complex computer vision problems. EX1083. A POSA with two years of experience would have also gained knowledge about machine learning, deep neural networks (DNNs), and simultaneous localization and mapping (SLAM), which involves edge detection and line identification. EX1083.

15. As these examples show, computer vision was an important topic during conferences sponsored by the automotive industry. The automotive industry was attending computer vision conferences. Companies focused on computer vision applications were founded and becoming integrated into the automotive industry. And graduates with Computer Science degrees were working on complex computer vision problems within the automotive industry. Accordingly, by October 4, 2021, an engineer with one- or two-years training and/or experience in design of software for camera/sensor systems in the automotive industry would have known about and have been familiar with 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, computerization of camera parameters from horizon line, etc.).

B. Motive's POSA definition adds additional skill to the POSA.

16. As I explained in the previous section, Motive's POSA definition would still include 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, computerization of camera parameters from horizon line, etc.) because someone with one- or two- years of experience in the design of software for camera/sensor systems would have had this knowledge. Motive's POSA definition also encompasses someone with a bachelor's in computer science, as my definition did. POR, 8. The main difference is therefore whether the POSA has experience in the automotive industry, as the chart shows below.

Definition in My Previous Declaration	Definition in Motive's POR
Someone with "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.). Additional education could serve as a substitute for the experience requirement." EX1003, ¶77.	"A POSITA for the '580 Patent as of 2021 would have a bachelor's degree in engineering or computer science with one- or two-years training and/or experience in design of software for camera/sensor systems in the automotive industry. " POR, 8-9 (emphasis added).

17. Shown above, Motive's POSA definition *adds* knowledge specific to

the automotive industry. I understand that it is generally easier to establish obviousness under a higher level of ordinary skill in the art. In other words, if the POSA has a higher level of skill, that POSA would have been more likely to be able to combine references with a reasonable likelihood of success.

18. Having additional knowledge about the automotive industry would have made combining computer vision techniques for autonomous driving applications more apparent. Before 2021, automotive engineers would have understood that perception—detecting lanes, vehicles, pedestrians, traffic signs, road conditions, etc.—would have been needed to enable driver assistance and autonomy, and that this perception must operate reliably in real-world driving environments. With this context, it would have been clear that computer vision techniques that accurately improve this perception were not only logical but necessary to build functional and scalable autonomous driving systems.

19. Accordingly, adding knowledge about the automotive industry, if anything, would have made the combinations more likely to be obvious.

III. I HAVE EXTENSIVE EXPERIENCE WITH COMPUTER VISION FOR THE AUTOMOTIVE INDUSTRY.

20. Motive argues that “unlike Mr. Andrews, Dr. Darrell reports no experience in the automotive industry or design of automotive sensor systems or any other factual basis to be knowledgeable about such ordinary knowledge or level of skill.” POR, 9-10. This could not be further from the truth.

21. I have extensive experience in the automotive industry, and especially with the design and function of automotive sensor systems for autonomous vehicles. I have had this expertise and experience since before the '580 patent's 2021 priority date, as demonstrated by my publications and work experience.

22. I have written numerous articles and academic papers on computer vision applications for autonomous vehicles, such as the following cited in my CV, which was submitted as Exhibit 1004:

- Xu, H., Gao, Y., Yu, F., & Darrell, T. (2017). End-to-End Learning of Driving Models from Large-Scale Video Datasets (pp. 2174-2182). Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. Retrieved from <https://arxiv.org/abs/1612.01079>
- Kim, J., Rohrbach, A., Darrell, T., Canny, J., & Akata, Z. (2018). Textual Explanations for Self-Driving Vehicles (pp. 563-578). Presented at the Proceeding of the European Conference on Computer Vision (ECCV). Referenced from http://openaccess.thecvf.com/content_ECCV_2018/papers/Jinkyu_Kim_Textual_Explanations_for_ECCV_2018_paper.pdf
- Wang, D., Devin, C., Cai, Q.-Z., Kr"ahenbu"hl, P., and Darrell, T. (2019). Monocular plan view networks for autonomous driving. In

2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 2876–2883. IEEE.

- Wang, D., Devin, C., Cai, Q.-Z., Yu, F., and Darrell, T. (2019). Deep object-centric policies for autonomous driving. In 2019 International Conference on Robotics and Automation (ICR), pages 8853–8859. IEEE.
- Kim, J., Moon, S., Rohrbach, A., Darrell, T., and Canny, J. (2020). Advisable learning for self-driving vehicles by internalizing observation-to-action rules. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9661–9670.
- Kim, J., Rohrbach, A., Akata, Z., Moon, S., Misu, T., Chen, Y.-T., Darrell, T., and Canny, J. (2021). Toward explainable and advisable model for self-driving cars. Applied AI Letters, 8pp, 2(4):e56.

23. I authored these six papers before the '580 patent's 2021 priority date.

24. Further, in May 2025, I authored a paper on a dataset and modeling framework for understructured roads for autonomous driving. EX1067. This paper built on the previous BDD100K dataset that was the largest open-source driving video dataset to help with autonomous driving development. EX1068; EX1069.

25. I also advised several PhD dissertations related to autonomous

driving. EX1070; EX1071. Each dissertation was completed before the '580 patent's 2021 priority date.

26. I have also worked with Nvidia and have spoken on deep learning breakthroughs in autonomous driving. EX1072; EX1073. Further, at the 2017 Berkeley EECS Annual Research Symposium, I spoke about autonomous driving and adapting from large-scale driving data. EX1074; EX1075.

27. Before 2021, I also participated in, gave invited talks at, and presented papers at the workshops on autonomous driving with CVPR and other conferences.

28. As discussed above in Section II.A, I participated in the 2016 IV where I met with people from Nexar. Nexar is ADAS dash cam company that has been operating since 2015. EX1076. After that meeting, I worked as the chief science advisor of Nexar from 2016 to 2024.

29. I am an affiliated director and was previously the faculty director of California's Partners for Advanced Transit and Highways (PATH). EX1077; EX1078. PATH collaborates with CalTrans to address challenges of California's surface transportation systems, including autonomous and non-autonomous driving. EX1079; EX1080. Under PATH, I am a founding co-director for UC Berkley DeepDrive (BDD), a multidisciplinary organization researching computer vision, robotics, and machine learning for automotive applications, such as autonomous driving. EX1081; EX1082. BDD collaborates with automotive

industry sponsors to develop new and emerging technologies with real-world applications in the automotive industry. EX1081; EX1082.

30. Therefore, as I show above, I have extensive experience in the automotive industry, and especially with the design and function of automotive sensor and camera systems for autonomous vehicles.

IV. I DISAGREE WITH MOTIVE’S IMPLICIT CONSTRUCTION OF “PREDICTIVE MODEL”

31. I disagree with Motive and Mr. Andrews that the ’580 patent’s predictive models exclude geometric approaches or geometric processes. Neither the specification nor the prosecution history supports such a narrow construction.

A. The ’580 patent does not exclude a geometric approach.

32. Motive and Mr. Andrews allege that the following statement in the ’580 patent redefines the term “predictive model” to exclude models that use a geometric approach to determine the horizon line. *See* EX1003, ¶69; POR, 5.

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

EX1001, 8:39-46.

33. The above excerpt does not define “predictive model” or exclude geometric calculations / a geometric approach. Rather, this excerpt states that a neural network can relax the requirement of needing three visible lane lines in frame. As the ’580 patent explains, the method would still work if only two lane lines were visible. *E.g.*, EX1001, 4:49-55 In such a case, missing lane lines would get predicted using geometry, specifically parallelism. EX1001, 4:49-55.

34. Rather than excluding geometric approaches, the ’580 patent encourages the use of them. The ’580 patent discloses that “[i]n other embodiments, a geometric process can be performed to identify lane and horizon lines.” EX1001, 8:12-14. In fact, in some embodiments where the image or video does not have high enough resolution, “the method can revert to a geometric algorithm to determine a lane width.” EX1001, 15:24-26.

35. A POSA would only have to look as far as the title of the patent to understand that a geometric process was not excluded from the invention. The title of the ’580 patent is: “Camera Initialization for Lane Detection and Distance Estimation *Using Single-View Geometry*.” EX1001, (54) (emphasis added).

36. For these reasons, the ’580 patent does not exclude the use of geometry to determine lane lines or horizon lines. The ’580 patent encourages it.

B. The prosecution history does not support a narrow construction of “predictive model.”

37. Mr. Andrews also alleges that the prosecution history of the ’580

patent supports his understanding that the “predictive model” in claim 1 excludes a geometric approach. EX1003, ¶¶25-28. I disagree.

38. Whereas Mr. Andrews analyzed only a portion of the prosecution history (EX1061, 135:11-21), my review shows that the Examiner allowed the claims not because of the predictive model limitation but because Kwant1 did not compute camera parameters, overlay lines on the video to generate an overlaid video, or transmit the overlaid video to an annotator device for manual review.

39. In the prosecution history, the Examiner issued only one Office Action. In that Office Action, the Examiner determined that the following portions of Kwant1 (EX1057) disclosed a method that “identif[ies] one or more lines in the video using a predictive model, the one or more lines including a horizon line.”

identifying one or more lines in the video using a predictive model *[the output of a lane feature detection process include lane models representing the detected lanes or features, and associated prediction confidence values indicating a probability that the processed features of the input image are predictive of the detected lane features]* 0046], the one or more lines including a horizon line *[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];

EX1002, 131.

40. The Examiner also determined that Kwant1 disclosed Steps A and C-

E. EX1002, 131-32. For Steps F and G, the Examiner relied on Kwant2. EX1002, 132-33.

41. In response, Motive argued that Kwant1 did not disclose using a predictive model to detect the horizon line.

In contrast to what is claimed, KWANT1 does not disclose using a <i>predictive model</i> to detect a horizon line. Instead, KWANT1 describes predicting <i>lane lines</i> . ¹ KWANT1 then describes extending these lane lines and, using a geometric approach, identifying the
horizon line. ² KWANT1 only describes using predictive modeling to predict lane lines. ³ Thus, KWANT1 does not describe predicting a horizon line as claimed.

EX1002, 145-46.

42. But this was not Motive’s only argument. Motive also raised four other arguments: Kwant1 did not describe a (1) “camera parameter,” (2) “overlaying” lines on a video, (3) “transmitting the overlaid video,” or (4) “receiving confirmation from the annotator device.” *See* EX1002, 146-148.

43. First, Motive argued that Kwant1 did not describe a “camera parameter” because Kwant1 only recites parametric coefficients of a lane line rather than a property of the camera itself. *See* EX1002, 146.

The Office Action cites to “12 parameters” in KWANT1 described in paragraph 113. However, these “parameters” are not parameters related to a *camera*, they are of a “parametric representation of a single lane line.”⁴ That is, these “parameters” define the parametric coefficients of a given lane line and thus define the lane line. These parameters simply have no relation to the camera used to generate the image but rather describe the lines—they are a pure mathematical representation of a line.

Applicant notes that the specification is clear that a camera parameter is a property of the camera itself (“(e.g., height, viewing angle, road plane normal, etc.)”, ¶ 5, and 6–7) and not a property of a lane line.

EX1002, 146.

44. Second, Motive argued that Kwant1 did not describe “overlaying” lines on a video. *See* EX1002, 147.

Applicant asserts that KWANT1 also fails to disclose “overlaying” lines on a video to generate an overlaid video. The Office Action cites to FIG. 4A of KWANT1 which depicts “polylines” placed over an image. However, this figure is not depicting the result of a video generation process. Rather, the figure is providing an explanation for what a polyline is via a graphical representation. Thus, while the figure in KWANT1 may arguably illustrate an image with polylines overlaid (a point not conceded), KWANT1 does not describe a method step for doing this (i.e., the explicit step of “overlaying” as claimed). Applicant asserts that KWANT1 must describe a *process* for generating an *image* and the failure to do so renders KWANT1 deficient.

EX1002, 147.

45. Third, Motive argued that Kwant1 did not describe “transmitting the overlaid video to an annotator device for manual review” because Kwant1 does not describe transmitting its image to anyone for manual review. *See* EX1002, 147-48.

Motive emphasized that Kwant1 does not describe “an annotator device.” *See* EX1002, 148. Motive also reiterated its argument that Kwant1 fails to describe overlaid lines. *See* EX1002, 148.

Applicant asserts that KWANT1 also fails to describe transmitting an overlaid video to an “annotator device for manual review.” *First*, KWANT1 never describes transmitting the image in FIG. 4A to anyone as it is a construction of the patent attorney or agent and not of a process, as discussed above. Further, the cited displays

(e.g., as described in paragraph 71) do not describe a video with overlaid lane or horizon lines. KWANT1 generally describes “using” a horizon to manipulate an image, but never actually “overlays” data on an image. Instead, for example, KWANT1 describes “localizing” a sign in an image based on the horizon⁵ or performing “sky detection” based on the horizon.⁶ At no point does KWANT1 describe displaying lines overlaid on an image. *Second*, even if the above operations were deemed a display of overlaid lines, the ultimate recipient in KWANT1 is a driver, not an “annotator device” as claimed.

EX1002, 147-48.

46. Finally, Motive argued that Kwant1 in view of Kwant2 did not describe “receiving confirmation from the annotator device” because, as they previously discussed, Kwant1 does not transmit its images for manual review and Kwant2’s learning module is not part of “manual review.” Instead, according to Motive, Kwant2 describes using “an automated machine learning technique” which Motive argued “is a fundamentally different process than using manual review to confirm the accuracy of a horizon detection.” EX1002, 148.

Applicant asserts that KWANT2 fails to describe a “confirmation” from an “annotator device.” The Office Action cites a “learner module” of KWANT2 as meeting this element. However, this learning module is not part of a “manual review.” Instead, it is an automated machine learning technique to adjust weights of a model based on a loss function. This is a fundamentally different process than using manual review to confirm the accuracy of a horizon detection.

EX1002, 148.

47. After the response, the Examiner issued a Notice of Allowance. In the Notice of Allowance, the Examiner maintained that Kwant1 discloses “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line.” *See* EX1002, 169-70; *see also* EX1061, 132:14-16, 133:2-5 (Mr. Andrews agreeing that the examiner maintained that Kwant disclosed a predictive model that identifies a horizon line). The Examiner found:

Prior art was found for the claims as follows:

Re. Claim 1, Kwant1 et al., [US Pub. No.: 2019/0034740 A1] discloses:

...

identifying one or more lines in the video using a predictive model *[the output of a lane feature detection process include lane models representing the detected lanes or features, and associated prediction confidence values indicating a probability that the processed features of the input image are predictive of the detected lane features]* 0046], the one or more lines including a horizon line *[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];

EX1002, 169-70.

48. The Examiner withdrew his position that Kwant1 disclosed:

- Step C: “computing at least one camera parameter based on the one or more lines;”
- Step D: “overlaying the one or more lines on the video to generate an overlaid video;” or
- Step E: “transmitting the overlaid video to an annotator device for manual review.”

See EX1002, 169-71.

49. Because the Examiner maintained that Kwant1 discloses “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line” but did not maintain its findings for Steps C-E, it is my opinion that the Examiner rejected Motive’s argument that the “predictive model” excludes geometric calculations. However, the Examiner found Motive’s arguments for Steps C, D, and E persuasive.

50. For these reasons, the prosecution history does not support Motive’s argument that the recited “predictive model” excludes a geometric approach.

V. GROUND 1: CHOE RENDERS OBVIOUS CLAIMS 1-4.

A. Choe discloses the “predictive model.”

51. I agree with Mr. Andrews’ initial opinion in his POPR declaration that Choe discloses predictive models that are used for sensor calibration to determine a set of calibration parameters. *See* POPR EX2003, ¶46. In his POR declaration, however, Mr. Andrews tries to walk away from his opinion, arguing that in Choe only algorithms—not predictive models—are used for sensor calibration to determine a set of calibration parameters. *See* EX2003, ¶46. As I will explain in this section, Mr. Andrews is now reading Choe too narrowly.

52. Choe states: “Based on driving statistics 123, machine learning engine 122 generates or trains a set of rules, algorithms, *and/or predictive models 124* for a variety of purposes.” EX1005, [0030] (emphasis added). When referring to item 124 later in the paragraph, Choe refers to “algorithms 124” and states “[a]lgorithms 124 may further include sensor calibration algorithms....” EX1005, [0030].

53. As the excerpt above states, Choe’s system can use rules, algorithms, and/or predictive models. EX1005, [0030]. So, when Choe refers to algorithms, a POSA would have understood that the algorithms could be used with predictive models for a particular purpose. For example, I agree with Mr. Andrews that an

algorithm can be used to implement a predictive model. EX1061, 70:6-12. Similar in Choe, an algorithm could be used to implement Choe's predictive models.

54. I also disagree with Motive and Mr. Andrews that using hardware parameters means that Choe does not use a predictive model to determine the horizon line. According to Motive and Mr. Andrews, Choe does not use predictive models to identify the horizon line because "Choe consistently describes throughout the specification that the horizon line is determined based on the hardware parameters of the camera." POR, 12 (citing EX2003, ¶43). However, Choe does not determine the horizon line based *only* on hardware settings. In Choe, the images are input into an image processor to determine the horizon line representing a vanishing point of a road. EX1005, [0045], Abstract. As a POSA would have understood, such a system could use both hardware settings and a predictive model to identify the horizon line accurately.

55. Motive and Mr. Andrews also argue that Choe's predictive models do not identify the horizon line because they are based on driving statistics. POR, 1; EX2003, ¶¶44-46. Motive states: "[A]ny 'predictive models' in Choe are based on the use of a machine learning module to process driving statistics, such as vehicle speeds, accelerations, etc., not based training a predictive model with images to identify lines in the images." POR, 13. Motive suggests that the '580 patent excludes the use of driving statistics in its predictive model. I disagree.

56. Choe describes:

Data collector 121 collects driving statistics 123 from a variety of vehicles, either autonomous vehicles or regular vehicles driven by human drivers. *Driving statistics 123 include information indicating the driving commands* (e.g., throttle, brake, steering commands) issued and *responses of the vehicles* (e.g., speeds, accelerations, decelerations, directions) captured by sensors of the vehicles at different points in time. Driving statistics 123 *may further include information describing the driving environments at different points in time*, such as, for example, routes (including starting and destination locations), MPOIs, road conditions, weather conditions, etc.

EX1005, [0029] (emphasis added).

57. Nowhere in the '580 patent is there any exclusion on the specific type of information that is used to train the predictive models. The '580 patent describes training its neural networks on annotated image frames and camera parameters. *See* EX1001, 7:5-8; 8:7-19. However, the '580 patent is clear that “[w]hile various embodiments have been described for purposes of this disclosure, such embodiments should not be deemed to limit the teaching of this disclosure to those embodiments.” EX1001, 19:61-64. Further, claim 1 merely recites “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line.” Claim 1 neither recites what exact inputs are used by the predictive model, nor does it explain the precise algorithm that performs the prediction.

58. As I explained in my previous declaration, Choe discloses a “predictive model.” EX1003, ¶¶90-97. Choe’s system predicts the initial horizon line. EX1005, [0048]. Then a user can provide input commands to move the horizon line to its correct location. EX1005, [0049].

59. To the extent the claims require a specific predictive model not disclosed in Choe, Choe in view of Khan/Ali (Grounds 2 and 5) and Choe in view of Workman (Grounds 3 and 6) both disclose specific neural networks that are used to identify the horizon line.

B. The calibration factors disclosed in Choe are or are used to compute “camera parameters.”

60. Motive argues that Choe does not compute a “camera parameter” because Choe only computes “calibration factors.” POR, 18. Motive’s argument is flawed because Choe’s “calibration factors” are camera parameters. Choe’s calibration factors are camera parameters because they relate to the camera itself and are associated with one or more operating characteristics of the camera.

61. After Choe computes calibration factors, Choe then uses the calibration factors to compute the camera’s pitch, yaw, and roll. For example, Choe describes “[a] first calibration factor or first correction value is determined for calibrating a pitch angle of the camera.” EX1005, [0049]. Choe also describes how a second calibration factor could similarly be used for a yaw angle of the camera. *See* EX1005, [0052]. Additionally, Choe describes “a third calibration

factor for calibrating a roll angle of the camera.” EX1005, [0053].

62. Pitch, yaw, and roll are also camera parameters because they relate to the camera itself and are associated with one or more operating characteristics of the camera. Mr. Andrews agrees that pitch, yaw, and roll are camera parameters. EX1061, 10:1-3.

63. I also disagree with Motive that Choe does not compute camera parameters based on the one or more lines (e.g., the horizon line and the lane lines). POR, 17-18. In Choe, the calibration factor for pitch, and therefore also the camera’s pitch angle, is determined based on the initial, predicted horizon line. EX1005, [0015], [0049], [0059], FIG. 7. In Choe, the calibration factors for yaw and roll, and therefore also the camera’s yaw and roll angles, are determined based on the initial, predicted lane lines. EX1005, [0017], [0052]-[0053].

VI. GROUND 2: CHOE IN VIEW OF KHAN/ALI RENDERS OBVIOUS CLAIMS 1-5.

64. For ground 2, I have been asked to respond to Mr. Andrews’ opinions on the motivation to combine and his argument that Khan/Ali does not identify the one or more lines using a predictive model. I address both in turn below.

A. My opinion on the motivation to combine Choe and Khan/Ali has not changed.

65. Mr. Andrews alleges that: “Absent the teachings of the ’580 patent, a POSITA would not be motivated to replace Choe’s hardware-based approach to

determine horizon lines with Khan/Ali geometric approach to identify a horizon line.” EX2003, ¶70. He argues this because Choe allegedly does not “motivate a POSITA to look for other references discussing the use of image-trained predictive models” in part because Choe allegedly does not “describe how ‘predictive models’ are used.” EX2003, ¶¶74, 90.

66. As I explained in my previous declaration, Choe uses predictive models for sensor calibration. EX1003, ¶¶93-94. Choe discloses a “machine learning engine 122 [that] generates or trains a set of rules, algorithms, and/or predictive models,” including for its sensor calibration algorithms. Choe, [0030]. A POSA before the 2021 priority date of the ’580 patent would have understood that Choe uses predictive models for sensor calibration. EX1003, ¶¶93-94. However, if a POSA was dissatisfied with Choe’s predictive models, the POSA would have searched for references that use predictive models and how such models were used. EX1003, ¶121. This would have led to Khan and Ali. EX1003, ¶121.

67. In fact, Mr. Andrews’ testimony appears to support this. Mr. Andrews determined that “Choe does not actually describe how ‘predictive models’ are used for any purpose.” EX2003, ¶90. Mr. Andrews also determined that Choe does not describe how you would determine a horizon line. EX1061, 77:13-24. According to Mr. Andrews, “[i]t would be nice if [Choe] provided more detail on that, but he doesn’t.” EX1061, 94:3-4. As I explained in my previous

declaration, if a POSA wanted that detail, the POSA would have found Khan/Ali, which explained how to use a predictive model to identify lane lines and a horizon line. EX1003, ¶121.

68. Further, as I discuss above, deep learning models were becoming increasingly important in the field of ADAS sensor design. I also explained why a POSA would have been motivated to use a predictive model for more accurate sensor calibration. EX1003, ¶¶123-27. Therefore, a POSA reading Choe would be motivated to seek out a predictive model such as the one described in Khan/Ali to fill in the gaps of Choe and use a more accurate sensor calibration method.

69. I was also not persuaded by Mr. Andrews' claim that Choe would not care about accuracy of the initial horizon line when calibrating its camera system. During deposition, Mr. Andrews alleged that "Choe doesn't care how accurate that candidate horizon line is." EX1061, 79:5-6. However, I find this statement to be inconsistent with basic design principles for autonomous vehicles.

70. A designer of autonomous driving systems would have wanted to have properly calibrated sensors to avoid accidents and harm to others. An improperly calibrated camera system for an autonomous vehicle could cause accidents with pedestrians or other vehicles. Therefore, a POSA would have understood that you would want to design a sensor system that gives accurate predictions of calibration elements such as the horizon line.

71. Further, even though manual calibration may be needed, a POSA would have understood that the accuracy of the initial, predicted horizon line is still important. Choe discusses how it calibrates its system in real time with manual adjustment. When Choe describes the current state of the art, it describes how “there has been a lack of efficient sensor calibration system[s].” Choe, [0003]. Choe describes how its algorithms and predictive models 124 can be used in real-time:

Algorithms 124 may further *include sensor calibration algorithms*, which will be described further below. Algorithms 124 can then be uploaded on ADVs *to be utilized during autonomous driving in real-time*.

EX1005, [0030] (emphasis added).

72. Choe further describes how, as part of its sensor calibration process, “a user can [] enter feedback” and then Choe’s system will adjust parameters so that the horizon line is properly identified. EX1005, [0046]-[0047]. A POSA would have understood that if Choe’s calibration is carried out in real-time, the initial predicted horizon line would need to be accurate for faster calibration. Mr. Andrews further supported this in his testimony, as shown below.

A Yeah, you could calibrate the sensors, I guess, in real time. That would be a rather risky thing to do *because if you don't have a good calibration to start with*, then your system is driving autonomously and it would be incorrect.

EX1061, 81:22-82:1 (emphasis added).

73. A POSA would have understood that having an accurate initial horizon line would not only result in more accurate calibration, it would also result in faster calibration. If the predicted horizon line is completely wrong, the reviewer of Choe's system would have to spend significant time manually correcting the horizon line. In real-time use, this would be cumbersome. Conversely, by accurately predicting the initial horizon line, the reviewer could quickly assess and correct minor inaccuracies. This would better enable real-time calibration.

74. For these reasons and the reasons in my prior declaration, which I will not repeat, a POSA would have been motivated to combine Choe and Khan/Ali.

75. If the Board adopts Motive's definition of a POSA, this definition would not change my opinions. A designer of automotive sensor system of ordinary skill would have been motivated to combine Choe's camera calibration system with Khan/Ali's predictive models.

B. Choe in view of Khan/Ali discloses the "predictive model."

76. Motive argues that "Khan and Ali disclose the use of a geometric algorithm, not a predictive model, to determine the horizon." POR, 24. I disagree.


77. As I explain above, the '580 patent does not exclude a geometric approach from the term “predictive model.” *See* Section IV. Under the plain and ordinary meaning of the term, Khan/Ali discloses a “predictive model.” EX1003, ¶¶133-39.

78. As I explained in my previous declaration and again in the following sections, Khan/Ali identifies one or more lines using a predictive model, the one or more lines including a horizon line. Khan/Ali discloses this element under three different interpretations of its algorithms. First, Khan/Ali uses a CNN to identify lane lines and the horizon line. Second, Khan/Ali’s horizon estimation algorithm is itself a predictive model that identifies the horizon line. Third, taken as a whole, the Khan/Ali’s end2end algorithm is a predictive model.

1. Khan and Ali’s CNN is a predictive model that is used to identify lane lines and the horizon line.

79. Khan and Ali use a neural network to determine lane lines. Khan explains: “[W]e proposed a novel lane detection algorithm that exploits the parallel nature of lane boundaries and fixed lane widths on the road to validate the detected lanes *and predict* lanes that are missed by the detector due to erased or effaced lane markings.” EX1007, 1824 (emphasis added). This “detection algorithm” includes lane prediction model which is a convolutional neural network (CNN) used to “predict [the] lane markers directly instead of sampling them from segmentation maps.” EX1007, 1825.


80. In Figure 3, Khan describes an “end2end” lane detection algorithm that includes the lane prediction model in the orange box (outlined in red below).



EX1007, Fig. 3 (annotated)

81. Khan/Ali describes how they “use a simple encoder-decoder based Convolutional Neural Network (CNN) to detect keypoints or markers [] along the lane boundaries.” EX1007, 1824. This CNN enables Khan/Ali’s system to detect lane boundaries even when lane markers are not present. EX1007, 1827-28. Khan describes how the CNN “take[s] motivation from keypoint detection in Mask-RCNN” (EX1007, 1825), which includes a backbone network and a network head. EX1007, 1826; EX1003, ¶133.

82. Figure 2 of Khan depicts an architecture diagram of the model:



EX1007, Fig. 2

83. This model uses an architecture with multiple convolution layers, batch normalization layers, a transpose convolution layer, ReLu (Rectified Linear Unit) layers, and pooling layers. *See* EX1007, 1826; EX1008, 1102. This model is a deep learning AI model trained to predict the lane lines based on identifying keypoints of the lane boundaries. *See* EX1008, 1102 (Khan uses “a fully convolutional deep network”).

84. As shown above, a POSA would have understood that the encoder network is a backbone, and the decoder network is the head. Specifically, Khan

uses the ResNet50 architecture as a backbone for encoding and then describes how its decoder network is a head for the overall network. *See* EX1007, 1826. The encoding network, for example the ResNet50 architecture, “encode[s] all relevant low and high level features from the input image.” EX1007, 1826. The decoder network then “learns the relationship between provided low and high level features and makes the decision on top of it.” EX1007, 1826.


85. Mr. Andrews agrees that the CNN in Khan’s Figure 2 includes a backbone and a head. Mr. Andrews agreed that the encoder network is a backbone. EX1061, 33:21-22 (“...it’s implied that they’re using an encoder network...as the backbone.”). He also agreed that the decoder network is a head. *See* EX1061, 33:23-34:4 (“Q: And in CNN, a person of ordinary skill in the art would have understood that like decoder is typically referred to as the head? A. Yes. So the Figure 2 model shows a backbone and a head? A: In a generic sense, yes.”).

86. The preferred embodiment of the ’580 patent also uses a CNN with a backbone and a head. The ’580 patent states that “[t]he disclosed embodiments utilize a convolutional neural network (CNN) to predict lane boundary and horizon markers....” EX1001, 4:44-46. The ’580 patent describes how “[a] backbone network generates a set of intermediate features representing the video.” EX1001, 1:55-57. “The intermediate features are then input into a lane or horizon prediction head (e.g., a neural network).” EX1001, 1:57-58.

87. Because Khan discloses a similar CNN as the '580 patent, I maintain my position that Khan/Ali discloses “identifying one or more lines in the video using a predictive model.”

88. Khan and Ali further disclose using this predictive model to identify a horizon line because Khan and Ali use the output of the CNN to determine the horizon line. The claims of the '580 patent recite “identifying one or more lines in the video *using* a predictive model, the one or more lines including a horizon line.” 20:5-7 (emphasis added). The claim language requires “using” a predictive model to identify one or more lines, including a horizon line. As I explain in more detail below, this is precisely what Khan/Ali does.

89. Khan/Ali identifies the horizon line based on the output of the CNN. As shown in Figure 3, after the CNN determines the lane boundaries, Khan describes how the algorithm determines the horizon line.



EX1007, Fig. 3 (annotated)

90. Khan describes how a “[c]ritical part of initialization is the estimation of a vanishing line (horizon) from these lines.” EX1007, 1826. Khan describes estimating the horizon lines based on the identified lane boundary keypoints. *See* EX1007, 1826 (“We start by fitting lines to the detected markers.”). Then, Khan discloses finding vanishing points to then use Ali’s horizon estimation algorithm. *See* EX1007, 1826-27 (“Once a set of forward and lateral VPs is found, we use [Ali’s] horizon estimation algorithm to find the best fitting horizon l_h .”).

91. Ali describes finding vanishing points based on the detected lane lines from the CNN. *See* EX1008, 1104 (“Since we already have detected lane lines in the image which lie on the road plane...”). Ali further discloses its algorithm which takes in the detected lane markers from the lane detector (i.e., the CNN) to output the estimated horizon line.

```
Algorithm 1: Best fitting Horizon ( $l_h$ ) estimation  
Input: Detected markers by lane detector  
Output: Estimated Horizon line  
1 Find forward and lateral VP for each frame in a video:  
2 initialize;  
3 for  $i \leq \text{num\_frames in video}$  do  
4   Fit lines to keypoints ( $B = \text{num\_lines}$ );  
5   if  $B \geq 2$  then  
6     for  $j \leq {}^B C_2$  do  
7        $\text{forward\_vpi}_j = \text{crossprod}(\text{line}_1, \text{line}_2)$   
8     end  
9   end  
10  if  $B \geq 3$  then  
11    for  $k \leq {}^B C_3$  do  
12      Get  $\text{lateral\_vpi}_k$  using cross ratios as described  
13      in Section 3.2.2  
14    end  
15  end  
16 RANSAC for best fitting horizon across video:  
17 initialize;  
18  $s = \text{random\_int between 1 to num\_forward\_vps};$   
19  $t = \text{random\_int between 1 to num\_lateral\_vps};$   
20 while  $\text{iterations} \leq t$  OR  $\text{num\_inliers} \leq \text{max\_inliers}$  do  
21   Fit line  $l_h$  to  $\text{forward\_vp}(s)$  and  $\text{lateral\_vp}(t)$ ;  
22   if  $\text{dist}(l_h, \text{forward\_vp}(s)) \leq \text{threshold}_s$  &  
23      $\text{dist}(l_h, \text{lateral\_vp}(t)) \leq \text{threshold}_t$  then  
24     record and count inlier points  
25   end  
26 end
```

EX1008, Algorithm 1: Best fitting Horizon (l_h) estimation

92. Therefore, Khan/Ali identifies a horizon line *using* a predictive model by estimating the horizon line using the outputs of the CNN. Mr. Andrews appears to agree. He testified as follows:

Q. To be clear, Algorithm 1 of Ali would use the CNN of Khan or the CNN of Khan plus the missing boundaries approach of Khan, right?

...

A. Yes.


EX1061, 44:16-22.

93. Based on Khan/Ali disclosing a CNN which is used to identify a horizon line, Choe in view of Khan/Ali discloses “identifying one or more lines in the video using a predictive model, the one or more lines including a horizon line.”

2. Khan/Ali’s horizon line estimation is a predictive model.

94. In addition to the estimation of the horizon line using the output of the CNN, Khan/Ali’s “horizon estimation” is also a predictive model.

95. Khan describes using Ali’s horizon estimation algorithm to find the horizon line. *See* EX1007, 1827 (“Once a set of forward and lateral VPs is found, we use [Ali’s] horizon estimation algorithm to find the best fitting horizon l_h .”). Ali describes that the horizon estimation model finds the best fitting horizon across the video using RANSAC. EX1008, 1104, Algorithm 1. RANSAC is a well-known machine learning model that is used as a tool for predictive modeling. The best fit horizon line is then overlaid onto the video as shown in Khan and Ali’s figures:



Left: EX1007, FIG. 1; Right: EX1008, FIG. 1


96. The '580 patent similarly determines markers (also called key points) for the lane and horizon lines and then “fits one or more lines to the markers.”

EX1001, 13:20-21 (“Next, in step 614, the method fits one or more lines to the markers.”). The '580 patent does not exclude algorithms from the scope of “predictive model.” Thus, a POSA would have understood that an algorithm that predicts the horizon line is a model that identifies a horizon line. Because Khan/Ali describes a prediction model that similarly estimates a best fit horizon line based on lane markers, a POSA would have understood that Khan/Ali’s horizon line estimation is a predictive model.

3. Khan and Ali's end2end algorithm for detecting lane lines and the horizon line is a predictive model.

97. In discussing claim 2, Mr. Andrews argues that “[a] POSITA would not consider the ‘end2end detection algorithm’ of Figure 3 in Khan (Ex. 1007), to be a predictive model.” EX2003, ¶95. I disagree.


98. Taken as a whole, the end2end lane detection algorithm described by Khan and Ali would be considered a “predictive model.” As I described above, Khan and Ali describe an algorithm that uses a CNN to detect lane markers, which then is used to estimate a horizon line. Khan and Ali's horizon line estimation could not happen without the use of the CNN.



EX1007, FIG. 3.

99. Mr. Andrews argues that the end2end algorithm is not a predictive model because it includes pre- and post-initialization steps. *See* EX2003, ¶¶95. However, a predictive model can have multiple steps.

100. In fact, the '580 patent itself describes its predictive model as having multiple steps to determine the lane lines, determine the horizon line, and initialize the camera. For example, “FIG. 3B is a block diagram illustrating a neural network for predicting lane and horizon lines and camera parameters according to some embodiments.” EX1001, 2:21-23. In Figure 3B, the predictive model is shown to have separate steps such as the Lane-Marking Head 306 and Camera initialization Head 310:



EX1001, FIG. 3B

101. The '580 patent further describes how the “backbone network 304 can convolve the input image 302 to generate a set of intermediate features.” EX1001, 7:17-19. This is shown as a separate step in Figure 4A to “process the image” before predicting the lane lines and horizon line. *See* EX1001, FIG. 4A, 10:5-9.




FIG. 4A

EX1001, FIG. 4A (annotated)

102. The '580 patent describes:

In step 404A, the method processes the image(s) using a backbone network. In an embodiment, the backbone network can comprise a backbone network such as backbone network 304, the disclosure of which is incorporated in its entirety.

EX1001, 10:5-9.

103. Additionally, the '580 patent describes post-processing steps to improve the prediction model based on a prediction loss. *See* EX1001, 10:21-49




FIG. 4A

EX1001, FIG. 4A (annotated)

104. In steps 408A and 410A, the '580 patent describes computing a prediction loss based on ground truth image labels to determine the performance of the predictive model. *See* EX1001, 10:21-44. If it is determined that there is a decrease in performance of the model or other stopping criteria is met, the method ends. *See* EX1001, 10:33-44. If the stopping criteria is not met, the method continues to step 412A, which adjusts the model by “comput[ing] partial derivatives against all trainable network parameters and back-propagates them to

adjust each layer parameter.” EX1001, 10:45-49.

105. The ’580 patent further describes how “the lane-marking head 306 can comprise a combined deep learning network that predicts both horizon and lane lines.” EX1001, 7:46-48. But, “[i]n another embodiment, lane-horizon marking head 306 can comprise two separate prediction heads” for the horizon and lane lines. EX1001, 7:48-50. Therefore, the ’580 patent explains that there can be separate steps to identify the horizon lines and the lane lines.

106. Based on the ’580 patent’s disclosure, there is nothing that suggests that the predictive model cannot have multiple steps.

4. Khan/Ali relaxes the requirement of having three lane lines in frame.

107. As I explained in Section IV.A, Motive alleges that the “predictive model” in the ’580 patent is different than the prior art because it relaxes previous requirements of having three lane lines in frame. The ’580 patent states:

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

EX1001, 8:39-46.

108. Khan/Ali relaxes this same requirement because Khan's approach detects missing lane boundaries. EX1007, 1827-28. For example, if there are only two lane lines visible in the image frame, Khan's model will identify the missing lane line. EX1007, 1827-28; EX1061, 39:25-40:6. As Mr. Andrews testified, the algorithm identifies the missing line so that the algorithm will have three lines, one of them is a synthetic line. EX1061, 37:8-13. The two lines visible on the image frame and the synthetically created third lane line would then get input into Ali's horizon estimation algorithm. EX1007, FIG. 3; EX1061, 40:12-41:15. Ali's horizon estimation algorithm will then determine the forward vanishing point and the lateral vanishing point to identify the best fitting horizon line across the video. EX1008, 1104, Algorithm 1; EX1061, 46:21-25.

109. Accordingly, Khan/Ali would determine the horizon line even if only two lane lines were visible. Mr. Andrews appears to agree, testifying:

Q. So Khan and Ali would determine the horizon line if only two lane lines were visible on the input image to the CNN predictive model, right?

A. They would determine a horizon line if only two lines were visible on the image and they synthetically derived a third line.

Q. And Khan describes how it synthetically derives a third line when that's the case, right?

A. Yes.

EX1061, 47:10-20.

110. Khan/Ali relaxes the requirement of needing three lane lines visible in the image frames, so Khan/Ali's predictive model also does not appear to differ from the one disclosed in the '580 patent on this basis. This is not overly surprising because the authors of Khan and Ali include the inventors of the '580 patent.

C. Choe and Khan/Ali perform calibration based on video.

111. To the extent Motive argues that Choe in view of Khan/Ali does not disclose "video," this is incorrect. Motive argues that Choe does not disclose receiving a video. POR, 16. However, Choe describes performing its calibrations on video, not just images. Further, to the extent Motive argues Choe does not perform its calibration on video, Ali affirmatively performs its processes on video. As I described in my previous declaration, Choe contemplates embodiments that include capturing video with video cameras. *See* EX1003, ¶86. Andrews further agreed with this point. *See* EX1061, 98:22-99:1.

112. In his deposition, Andrews said:

Q So Choe contemplates embodiments that use video cameras, right?

A It's still cameras and/or video cameras, yes.

Q And video cameras capture video, right?

A Yes.

Q And a video is a set of image frames?

A A video is a set of image frames.

EX1061, 98:22-99:3.

113. To the extent that Motive still argues that Choe does not disclose a video, Khan/Ali performs its algorithm on a video. Ali describes performing its vanishing point estimation across frames of a video. EX1008, 1104. Andrews also agreed that Ali's algorithm is performed on video. *See* EX1061, 49:1-4, 52:20-53:1.

Q So in Ali's Algorithm 1, it's fitting a horizon line to the forward and lateral vanishing points of the image frames within a video, right?

A That's my understanding, yes.

EX1061, 49:1-4.

Q Right. And the example that Khan and Ali presented in their papers, the estimation of horizon and IPM is on a short video clip. Do you see that?

A Yes.

Q Do you have any reason to dispute the accuracy of that statement?

A No.

EX1061, 52:20-53:1.

114. Therefore, Choe in view of Khan/Ali teaches identifying one or more lines from image frames of a “video.”

VII. GROUND 3: MY OPINION ON THE MOTIVATION TO COMBINE CHOE AND WORKMAN HAS NOT CHANGED.

115. Similar to his argument for Ground 2, Mr. Andrews argues that a POSA would not have been motivated to combine Workman’s predictive model with Choe’s calibration system because Choe uses a “camera-hardware-based approach to determine initial horizon lines” instead of predictive models. *See* EX2003, ¶¶73-74. As I explain throughout my previous declaration and this declaration, I disagree with the premise that Choe does not use predictive models for sensor calibration. Even Mr. Andrews himself interpreted Choe as disclosing predictive models for sensor calibration but has since changed that opinion. *Compare* POPR EX2003, ¶46, *with* EX2003, ¶46.

116. As I explained in my previous declaration, Choe uses predictive models for sensor calibration. Choe discloses a “machine learning engine 122 [that] generates or trains a set of rules, algorithms, and/or predictive models,” including for its sensor calibration algorithms. Choe, [0030]. A POSA before the 2021 priority date of the ’580 patent would have understood that Choe uses predictive models for sensor calibration. However, if a POSA was dissatisfied with Choe’s predictive models, the POSA would have been motivated to search for references that use predictive models and how such models were used. EX1003,

¶146. This would have led to Workman, which provides state-of-the-art predictive models for identifying the horizon line. EX1003, ¶¶146-50.

117. In fact, Mr. Andrews' testimony appears to support this. Mr. Andrews determined that "Choe does not actually describe how 'predictive models' are used for any purpose." EX2003, ¶90. Mr. Andrews also determined that Choe does not describe how you would determine a horizon line. EX1061, 77:13-24. According to Mr. Andrews, "[i]t would be nice if [Choe] provided more detail on that, but he doesn't." EX1061, 94:3-4. As I explained in my previous declaration, if a POSA wanted that detail, the POSA would have found Workman, which explained how to use a predictive model to identify the horizon line. EX1003, ¶146.

118. Mr. Andrews also alleges that a POSA would not have combined Choe and Workman because "Workman is a generic computer vision reference that has no relation to the type of automotive system of Choe." EX2003, ¶72. I disagree. A POSA would have understood that Workman was relevant to autonomous driving vehicles (ADVs), such as the one disclosed in Choe.

119. Computer vision is a broad field that relies on properly calibrated cameras. The problems that general computer vision systems experience with initializing cameras and other sensors are similar to the problems that occur with initializing cameras in the specific computer vision application of autonomous driving. Therefore, in my opinion, I disagree with Motive's claim that the problem

of camera initialization in an automotive application is “different than the type of problems encountered in computer vision generally.” POR, 8. A designer with two years of experience in the automotive industry would have known that methods used in general computer vision would be applicable to automotive systems. A POSA would understand that, to initialize a camera for use in an autonomous automotive application, various predictive modeling algorithms could be used, including the CNN as described in Workman.

120. In the opening sentence of Workman, Workman describes how “[s]ingle image horizon line estimation is one of the most fundamental geometric problems in computer vision.” EX1011, 0001. Workman then describes how “[k]nowledge of the horizon line enables a wide variety of applications, including: ... geometrically biased pedestrian and vehicle detection.” EX1011, 0001. As a POSA would have understood, pedestrian and vehicle detection is directly pertinent to ADVs, such as the one Choe disclosed.

121. Workman's figures also show its application to vehicles:



EX1011, 0009.

122. Though Workman is not directed *solely* to automotive systems, Workman was a state-of-the-art predictive model that outperformed existing models for identifying the horizon line. EX1003, ¶¶147-49. As I explained above in Section II.A, a designer for autonomous driving sensor systems would have been aware of the intersection of computer vision and automotive systems. A POSA would have also understood that general computer vision advancements were readily applicable to sensor systems for ADVs. Therefore, a POSA under either definition would have found Workman highly relevant to Choe's ADV.

123. Additionally, Mr. Andrews argues that Workman's faster performance "is not relevant to Choe." EX2003, ¶78. Mr. Andrews further argues that "[t]he

fact that the experiments in Workman required ‘only milliseconds per image’ and ‘outperform[ed] all existing methods’ does not provide any indication of how it performs compared to the camera-hardware-based approach of Choe.” EX2003, ¶78. I disagree.

124. Performance is incredibly important for camera calibration, especially in an ADV like Choe that is calibrated in real time. *See* EX1005, [0017], [0054]. Above, I described how Mobileye, with its specialized processing chips for ADAS, was revolutionary. Its chips and products were the reason it was purchased by Intel for billions of dollars in 2017. EX1064. Fast processing of images in autonomous driving is critical to having better and more responsive systems. Therefore, a POSA would be heavily motivated to seek out the best methods and systems for processing images.

125. As I discuss above, Choe teaches performing its horizon prediction on video frames. A POSA would have understood that Workman’s image processing could be applied to Choe’s video frames. Therefore, a POSA would have understood that Choe in view of Workman would perform its horizon prediction on video frames, not just images.

126. Choe describes how camera calibration is performed “while the vehicle is driving and real-time images are captured.” EX1005, [0017], [0054]. Workman’s ability to process images to determine the horizon line in milliseconds,

several orders of magnitude faster than current state-of-the-art methods, would have motivated a POSA to integrate it into Choe's camera calibration. EX1003, ¶148. Mr. Andrews' testimony to the contrary does not show otherwise.

127. Similar to Ground 2, I will not repeat the remarks from my previous declaration on the full motivation to combine Choe and Workman. I stand by my previous explanation of the motivation to combine Choe and Workman.

128. If the Board adopts Motive's definition of a POSA, as I explain above, this definition would not change my opinions. A designer of automotive sensor system of ordinary skill would have been motivated to combine Choe's camera calibration with Workman's predictive model.

VIII. GROUND 4: MY OPINION ON THE MOTIVATION TO COMBINE CHOE IN VIEW OF WESTMACOT HAS NOT CHANGED.

129. For ground 4, Motive alleges that "Westmacot teaches away for doing any camera calibration" and therefore a POSA would not be motivated to compute and calibrate the camera height and the road plane normal, as taught in Westmacot. *See* POR, 56. I disagree.

130. As an initial matter, there are many different types of "calibration" for cameras. For example, there is calibration of the intrinsic parameters of the camera and calibration of the extrinsic parameters of the camera. "Intrinsic" parameters of a camera include the relationships between pixel and camera coordinates, focal length of the camera, and other intrinsic characteristics of the camera. *See* EX1084,

Slide 4. “Extrinsic” parameters of a camera include the camera’s location and orientation in the physical world, such as the camera’s pose relative to the road plane normal. *See* EX1084, Slide 4. Because the ’580 patent talks about determining the orientation of the camera and the camera height, the ’580 patent is talking about calibrating the “extrinsic” parameters of the camera.

131. The portion of Westmacot that Motive cites does not teach away from extrinsic calibration. The “calibration” Westmacot is talking about is linked to complex “calibrated cameras and Lidar,” as compared to “uncalibrated equipment, such as a mobile phone or a dash-cam.” EX1010, 19:5-6, 37:10-13. This is related to calibration of the intrinsic parameters of these types of cameras. Throughout its disclosure, Westmacot describes whether an image capture device is “calibrated” as indicating whether the image capture device is complex such as Lidar as compared to “uncalibrated” cameras, like a dash cam. *See* EX1010, 19:5-6, 37:10-13, 46:27-29. Westmacot’s purported improvement over the art is its ability to annotate images and video using “uncalibrated cameras” because the camera parameters and road structure can be “determined using only information from image data captured by the image capture device.” EX1010, 3:26-31. “For example, one or more parameters of the image capture device ... may be derived from (only) the image data captured by the image capture device.” EX1010, 3:31-34. These parameters include the camera height and the road plane normal, or

extrinsic camera parameters. *See* EX1010, 32:29-33:2 (“These reference parameters comprise the following image capture parameters, relating to the position and orientation of the image capture device 302 of the training vehicle 300: ... height H of the camera...horizontal offset S of the camera...the orientation of the camera 302 relative to the vehicle”).

132. Therefore, a POSA would have understood that Westmacot is describing that it is desirable to have a system for calibration of “one or more parameters of the ...device...[or]...road structure from....the [uncalibrated, low cost] image capture device.” *See* EX1010, 3:31-34. Therefore, Westmacot is describing how it would be preferable to use a low-cost camera for the purpose of collecting data that could estimate the camera and/or scene parameters.

133. However, to the extent Motive argues that the “calibration” Westmacot is referencing is whether one should calibrate a camera for autonomous driving at all, Westmacot is not teaching away from all calibration of a camera. Westmacot merely expresses a preference for not requiring calibration, stating: “*Preferably* the expected road structure is determined ... without any requiring any calibration of the image capture device.” EX1010, 3:25-31 (emphasis added). I understand that a preference is not enough to constitute a teaching away. In an ideal world, a camera does not need to have its intrinsic and extrinsic parameters calibrated. However, in the real-world, all parameters of a sensor need to be

calibrated and re-calibrated. For example, Choe explains that sensors require periodic calibration to maintain the accuracy of the sensors so “that the perception system of an autonomous driving system can accurately detect and recognize the objects and lane configuration of the driving environment.” EX1005, [0003].

Therefore, though Westmacot has a preference not to require calibration, a POSA would still have understood that in real-world conditions, calibration would be necessary. Indeed, Choe requires calibration and does not teach away from calibrating its camera. Nor does Choe teach away from computing the camera height and the road plane normal.

134. Westmacot’s preference is not, as Mr. Andrews argues, “the opposite of what Choe teaches.” EX2003, ¶101. Mr. Andrews cites a passage of Choe that describes how “an image processed by perception module 302 may be utilized by sensor calibration module 308 for sensor calibration.” EX2003, ¶101 (quoting EX1005, [0046]). This passage does not show how Choe is the opposite of Westmacot. Rather, this passage describes how Choe uses images for sensor calibration, similar to how Westmacot’s image labeling system can be used to calibrate sensors.

135. Additionally, Westmacot’s ability to use only image data to annotate video frames with lane and horizon lines would further motivate a POSA to combine Westmacot and Choe. Westmacot describes how using image processing

instead of conventional calibration techniques would enable the system to use “inexpensive and un-calibrated equipment, such as a mobile phone or dash-cam costing a few hundred pounds, or other low cost consumer equipment.” EX1010, 37:10-13. As described above with respect to Workman, a designer of automotive sensor systems would have wanted to design a system that was efficient and cost effective. Therefore, a POSA would have been motivated to calculate the camera height and road plane normal through image processing, as Westmacot teaches, to allow Choe’s system to use inexpensive cameras.

136. Motive argues that Choe in view of Westmacot does not teach “computing at least one camera parameter based on the one or more lines.” *See* POR, 57. Motive and Mr. Andrews claim that Westmacot’s disclosure of deriving lane lines from equations that use the camera height and road plane normal has no application to Choe because “Choe does not reconstruct the path of the image capture device.” *See* POR, 57-58; EX2003, ¶102. I disagree.

137. In my previous declaration, I discussed how Choe describes calculating lane lines and how Westmacot discloses specific equations for determining those lane lines. EX1003, ¶¶160-61. Westmacot further describes that one may want to calculate the rotational orientation of the camera based on the road plane normal. EX1003, ¶162; EX1010, 9:19-22. As I previously discussed, a POSA would have been motivated to modify Choe with Westmacot to calculate

the road plane normal because it is necessary to determine the rotational orientation of the camera. EX1003, ¶162. Using Westmacot's equation and determined lane lines, a POSA can modify Choe to calculate the road plane normal. For similar reasons, a POSA would have been motivated to modify Choe with Westmacot's calculation of the camera height.

138. As previously discussed, Westmacot teaches how to calculate camera parameters from image data. As discussed in my previous declaration, Choe does not explicitly disclose how one would calculate the camera parameters. *See* EX1003, ¶160. Therefore, a POSA would have been motivated to look at Westmacot for a specific way to calculate a camera parameter.

139. Similar to Grounds 2 and 3, I will not repeat the remarks from my previous declaration with respect to further reasons a POSA would have been motivated to combine Choe and Westmacot. I stand by my previous explanation of the motivation to combine Choe and Westmacot.

140. If the Board adopts the Motive's definition of a POSA, as I explain above, this definition would not change my opinions. A designer of automotive sensor system of ordinary skill would have been motivated to combine Choe and Westmacot.

IX. GROUNDS 5 AND 6: MY OPINION ON THE MOTIVATION TO COMBINE HAS NOT CHANGED

141. For ground 5, Motive alleges that “Petitioner fails to provide any motivation why a POSITA would combine Choe, Westmacot, Khan, and Ali all together.” POR, 58. For ground 6, Motive alleges that “Petitioner fails to provide any motivation of why a POSITA would combine Choe with both Westmacot and Workman.” POR, 59. I disagree.

142. As I explained in my previous declaration, a POSA would have been motivated to combine Choe and Westmacot. EX1003, ¶¶159-74. A POSA would have also been motivated to combine Choe and Khan/Ali. EX1003, ¶¶121-31. These rationales are independent and relate to different parts of Choe’s system. If a POSA were to make both combinations, the result would therefore disclose every element of the challenged claims.

143. I also explained in my previous declaration that a POSA would have been motivated to combine Choe and Westmacot. EX1003, ¶¶159-74. A POSA would have also been motivated to combine Choe and Workman. EX1003, ¶¶145-50. These rationales are independent and relate to different parts of Choe’s system. If a POSA were to make both combinations, the result would therefore disclose every element of the challenged claims.

X. CONCLUSION

144. In signing this declaration, I recognize that the declaration will be filed as evidence in an *inter partes* review before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required, I will appear for cross-examination within the United States during the time allotted.

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

Executed on this 11th day of February 2026, in Berkeley, CA.



Trevor J. Darrell, Ph.D.

Date: 2/11/2026