

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

TESLA, INC.,

Petitioner,

v.

PERCEPTIVE AUTOMATA LLC,

Patent Owner.

Case No. IPR2025-01575

Patent No. 11,753,046

PATENT OWNER'S RESPONSE

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	THE '046 PATENT AND THE CHALLENGED CLAIMS	3
	A. Overview of the '046 Patent.....	3
	B. Claim Overview	4
III.	LEVEL OF ORDINARY SKILL	7
IV.	CLAIM CONSTRUCTION	8
V.	APPLICABLE LEGAL STANDARDS.....	10
	A. Tesla's Burden.....	10
	B. Obviousness in IPR Proceedings.....	11
	C. The <i>Graham</i> Factors in IPR Proceedings.....	12
VI.	THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 1	13
	A. Cox	14
	B. Cox Does Not Teach “user”/“users”	16
	C. Cox Does Not Teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers”	16
	D. Cox Does Not Teach “storing the summary statistics in association with the image as part of the training data”	21
	E. Cox Does Not Teach “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image”	23
	F. Challenged Dependent Claims	26

G.	Tesla and its Expert Fail to Analyze the <i>Graham</i> Factors, As Required for a Proper Obviousness Analysis	26
VII.	THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 2	27
A.	Ross	28
B.	Cox and Ross Do Not Teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers”	28
C.	Cox and Ross Do Not Teach “storing the summary statistics in association with the image as part of the training data”	29
D.	Cox and Ross Do Not Teach “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image”	29
E.	Challenged Dependent Claims	30
F.	A POSITA Would Not Have Been Motivated to Combine Cox and Ross in the Manner Tesla Alleges.....	31
VIII.	THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 3	35
IX.	THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 4	35
A.	Munro	36
B.	Cox and Munro Do Not Teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers”	36
C.	Cox and Munro Do Not Teach “storing the summary statistics in association with the image as part of the training data”	38
D.	Cox and Munro Do Not Teach “the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image”	40

E.	Challenged Dependent Claims	41
F.	A POSITA Would Not Have Been Motivated to Combine Cox and Munro in the Manner Tesla Alleges.....	42
G.	The Petition Fails to Analyze the <i>Graham</i> Factors, As Required for a Proper Obviousness Analysis.....	43
X.	THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 5	44
A.	Ellenbogen.....	44
B.	Ellenbogen and Munro Do Not Teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers”	45
C.	Ellenbogen and Munro Do Not Teach “storing the summary statistics in association with the image as part of the training data”	47
D.	Ellenbogen and Munro Do Not Teach “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image”	48
E.	Challenged Dependent Claims	50
F.	A POSITA Would Not Have Been Motivated to Combine Ellenbogen and Munro in the Manner Tesla Alleges.....	50
XI.	TESLA FAILED TO SHOW THAT CERTAIN LIMITATIONS ARE PRINTED MATTER NOT ENTITLED TO PATENTABLE WEIGHT	52
A.	The Claim Limitations At Issue Are Not Printed Matter.....	53
B.	The Alleged Printed Matter Is Functionally Related to the Substrate	57
XII.	Conclusion	60

TABLE OF AUTHORITIES

Cases

<i>Amazon Web Services, Inc. v. Saint Regis Mohawk Tribe</i> , IPR2019-00103, Paper 22 (P.T.A.B. May 10, 2019)	10
<i>Arendi S.A.R.L. v. Apple Inc.</i> , 832 F.3d 1355 (Fed. Cir. 2016)	26
<i>C R Bard Inc. v. Angiodynamics, Inc.</i> , 979 F.3d 1372 (Fed. Cir. 2020)	53, 57
<i>CFMT, Inc. v. YieldUp Int’l Corp.</i> , 349 F.3d 1333 (Fed. Cir. 2003)	11
<i>eBay, Inc. v. Paid, Inc.</i> , CBM2014-00125, Paper 15 (P.T.A.B. Sept. 30, 2014).....	13, 27, 43
<i>Eizo Corp. v. Barco N.V.</i> , No. IPR2014-00358, Paper 11 (P.T.A.B. July 23, 2014).....	13, 27, 43
<i>FanDuel, Inc. v. Interactive Games LLC</i> , 966 F.3d 1334 (Fed. Cir. 2020)	52
<i>Graham v. John Deere Co.</i> , 383 U.S. 1 (1966).....	<i>passim</i>
<i>Harmonic Inc. v. Avid Tech., Inc.</i> , 815 F.3d 1356 (Fed. Cir. 2016)	<i>passim</i>
<i>In re DiStefano</i> , 808 F.3d 845 (Fed. Cir. 2015)	53, 54, 55, 58
<i>In re Ethicon, Inc.</i> , 844 F.3d 1344 (Fed. Cir. 2017)	17, 18, 19, 20
<i>In re Gulack</i> , 703 F.2d 1381 (Fed. Cir. 1983)	52, 58
<i>In re Kahn</i> , 441 F. 3d 977 (Fed. Cir. 2006)	12

<i>In re Lowry</i> , 32 F.3d 1579 (Fed. Cir. 1994)	<i>passim</i>
<i>In re McKee</i> , 64 F.2d 379 (CCPA 1933)	54
<i>In re Miller</i> , 418 F.2d 1392 (CCPA 1969)	57, 58
<i>In re NuVasive, Inc.</i> , 841 F.3d 966 (Fed. Cir. 2016)	24, 27
<i>In re Rijckaert</i> , 9 F.3d 1531 (Fed. Cir. 1993)	11
<i>In re Royka</i> , 490 F.2d 981 (C.C.P.A. 1974)	11
<i>Juniper Networks, Inc. v. Swarm Tech. LLC</i> , IPR2021-01445, Paper 62 (P.T.A.B. Feb. 7, 2023)	24, 27
<i>King Pharms., Inc. v. Eon Labs., Inc.</i> , 616 F.3d 1267 (Fed. Cir. 2010)	54
<i>Kyocera Senco Indus. Tools Inc. v. Int’l Trade Comm’n</i> , 22 F.4th 1369 (Fed. Cir. 2022)	7
<i>Liberty Mut. Ins. Co. v. Progressive Cas. Ins. Co.</i> , CBM2012-00003, Paper 8 (PTAB Oct. 25, 2012)	<i>passim</i>
<i>Moses Lake Indus., Inc. v. Enthone, Inc.</i> , IPR2014-00246, Paper 6 (P.T.A.B. June 18, 2014)	13, 27, 43
<i>Praxair Distrib., Inc. v. Mallinckrodt Hosp. Prods. IP Ltd.</i> , 890 F.3d 1024 (Fed. Cir. 2018)	57
<i>TQ Delta, LLC v. Cisco Sys.</i> , 942 F.3d 1352 (Fed. Cir. 2019)	<i>passim</i>
<i>Wellman, Inc. v. Eastman Chem. Co.</i> , 642 F.3d 1355 (Fed. Cir. 2011)	10

<i>Xerox Corp. v. Bytemark, Inc.</i> , IPR2022-00624, Paper 12 (P.T.A.B. Feb. 10, 2023)	17, 18, 19, 20
--	----------------

Statutes

35 U.S.C. § 103(a)	11
35 U.S.C. § 316(e)	52

Other Authorities

OFFICE PATENT TRIAL PRACTICE GUIDE.....	12
---	----

Regulations

37 C.F.R. § 42.120	1
37 C.F.R. § 42.22(a)(2)	11, 38, 39, 41
37 C.F.R. § 42.104(b)(4).....	11, 38, 39, 41

PATENT OWNER’S EXHIBIT LIST

Exhibit #	Description
2001	Tesla February 1, 2022 IDS
2002	Tesla March 1, 2022 IDS
2003	Tesla March 4, 2022 IDS
2004	Tesla March 17, 2022 IDS
2005	Tesla March 18, 2022 IDS
2006	Tesla March 18, 2022 IDS
2007	Tesla March 21, 2022 IDS
2008	Tesla March 22, 2022 IDS
2009	Tesla March 22, 2022 IDS
2010	Tesla May 4, 2022 IDS
2011	Tesla May 10, 2022 IDS
2012	Tesla May 11, 2022 IDS
2013	Tesla May 11, 2022 IDS
2014	Tesla May 11, 2022 IDS
2015	Tesla May 12, 2022 IDS
2016	Tesla May 12, 2022 IDS
2017	Tesla May 13, 2022 IDS
2018	Tesla May 13, 2022 IDS
2019	Tesla May 25, 2022 IDS
2020	Tesla June 28, 2022 IDS
2021	Perceptive’s Original Complaint
2022	Tesla’s Motion to Dismiss
2023	Perceptive’s First Amended Complaint
2024	Tesla’s Motion to Stay
2025	Perceptive’s Opposition to Tesla’s Motion to Stay
2026	Tesla’s Reply in Support of Motion to Stay
2027	Perceptive’s Sur-Reply in Opposition to Motion to Stay
2028	Docket Sheet
2029	Motion for Protective Order
2030	Protective Order
2031	Operative Docket Control Order
2032	Perceptive’s Document Request Letter
2033	Perceptive’s First Set of Interrogatories
2034	Tesla’s First Set of Interrogatories
2035	Tesla’s Notice of Broadened Sotera Plus Stipulation

2036	Tesla's Initial and Additional Disclosures
2037	District Court Statistics on Post-Institution Motions to Stay Pending IPR
2038	District Court Time-to-Trial Statistics
2039	Declaration of Eli Saber, Ph.D.
2040	CV of Eli Saber, Ph.D.
2041	Lord, F.M. (1951), A THEORY OF TEST SCORES AND THEIR RELATION TO THE TRAIT MEASURED. ETS Research Bulletin Series, 1951: i-126.
2042	Kooij JFP, Schneider N, Flohr F, Gavrilă DM (2014) Context-based pedestrian path prediction. In: Proceedings of the European Conference on Computer Vision (ECCV). Berlin: Springer, pp. 618–633
2043	A. Zyner, S. Worrall, J. Ward and E. Nebot, "Long short term memory for driver intent prediction," 2017 IEEE Intelligent Vehicles Symposium (IV), Los Angeles, CA, USA, 2017, pp. 1484-1489
2044	C. Aggarwal, "Neural Networks and Deep Learning", Springer ISBN 978-3-319-94462-3
2045	Deposition of Dr. Jason Janét

I. INTRODUCTION

Pursuant to 37 C.F.R. § 42.120, Patent Owner Perceptive Automata LLC (“Perceptive”) submits this Response to the Petition for *Inter Partes* Review (“Petition”) filed by Tesla, Inc. (“Tesla”) challenging claims 1-19 of U.S. Patent No. 11,753,046 (the “’046 patent”) (Ex. 1001).

The Petition presents five grounds of obviousness, one based on a single reference (International Publication No. WO2014/210334 (“Cox”) (Ex. 1007)), three primarily based on Cox, and one based on U.S. Patent Application Publication No. 2017/0099200 (“Ellenbogen”) (Ex. 1004) and U.S. Patent Application Publication No. 2016/0162456 (“Munro”) (Ex. 1006).

Cox discloses processes whose aim is to improve machine learning models for recognizing faces in images. In contrast, the ’046 patent claims concern processes and systems for a different application of machine learning—specifically, predicting what road users (motorists, bicyclists, pedestrians, etc.) will do when navigating streets with an autonomous vehicle. As discussed below, Cox does not teach several elements of the challenged claims—most notably the requirements of generating and predicting “summary statistics describing the state of mind of the user in [an] image”—which is significant for the ’046 patent’s autonomous driving applications but not for face detection. While Tesla and its expert (inconsistently) point to myriad features of Cox as allegedly being summary statistics, each theory fails, which is

fatal to all of Tesla's Cox-based obviousness grounds. Tesla's Cox-based grounds have additional shortcomings, including failure to properly consider the *Graham* factors and lack of motivation to combine Cox with the other references.

Tesla's Ellenbogen-Munro ground has similar deficiencies. Munro concerns natural language processing of documents, not any type of machine learning relating to images (for autonomous driving or otherwise). Despite Munro's focus on a different field, Tesla incorrectly relies on its response aggregation process as supposedly teaching "summary statistics describing the state of mind of the user in [an] image"—an argument just as weak as it sounds on its face. Munro's different focus also means that a POSITA would not have been motivated to combine it with Ellenbogen.

Finally, Tesla attempts to argue that many limitations of the challenged claims should not be given patentable weight in the invalidity analysis under the printed matter doctrine. Tesla's arguments come woefully short on both the law and the facts and should be rejected.

Given these flaws in Tesla's theories, the Board should conclude that challenged claims 1-19 are not unpatentable.

II. THE '046 PATENT AND THE CHALLENGED CLAIMS

A. Overview of the '046 Patent

The inventions disclosed and claimed in the '046 patent generally relate to systems, computer-implemented methods, and computer readable media associated with controlling autonomous vehicles by “judging the behavior of people near a vehicle in a way that more closely resembles how human drivers would judge the behavior.” Ex. 1001 at 3:58-60. The '046 patent discloses an implementation for achieving this result by obtaining responses to images displaying road scenes to determine a road user’s “state of mind—e.g. that a pedestrian is likely to cross the street or that an oncoming vehicle is unlikely to be willing to yield to the vehicle containing the sensor if the vehicle containing the sensor needs to turn.” *Id.* at 10:34-53. The images of road users (e.g., motorists, cyclists, pedestrians) in a road scene can be captured by the autonomous vehicle and sent to a remote location for obtaining responses and training a model. *Id.* at 7:56-8:14. Training dataset(s) comprising summary statistics of the responses by human labelers are used to train a learning based model so that the learning based model can predict whether a road user in a new image has a given “state of mind.” *Id.* at 10:34-53. “The trained model or algorithm makes a prediction of what a pedestrian or other person shown in the ‘live data’ would do based on the summary statistics and/or training labels” *Id.* at 7:32-35. “The predictions produced by the trained model comprise a set of

predictions of the state of mind of road users that can then be used to improve the performance of autonomous vehicles, robots, virtual agents, trucks, bicycles, or other systems that operate on roadways by allowing them to make judgments about the future behavior of road users based on their state of mind.” *Id.* at 7:32-35.

B. Claim Overview

Independent claims 1, 8, and 15 are reproduced below:

Claim 1	Limitations
1[pre]	A computer-implemented method comprising:
1[a]	storing a plurality of images, each image displaying one or more users;
1[b]	generating training data from the plurality of images, the generating comprising, for each image:
1[b-i]	sending the image to a plurality of human observers, each human observer presented with a request to answer a question about a state of mind of a user in the image,
1[b-ii]	receiving, from each of the plurality of human observers, a response representing a judgment by the human observer of the state of mind of the user in the image,
1[b-iii]	generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers, and
1[b-iv]	storing the summary statistics in association with the image as part of the training data;

1[c]	training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image; and
1[d]	executing the trained model to predict a state of mind of a user in a new image.
Claim 8	Limitations
8[pre]	A non-transitory computer readable storage medium storing instructions that when executed by one or more processors, cause the one or more processors to perform steps comprising:
8[a]	storing a plurality of images, each image displaying one or more users;
8[b]	generating training data from the plurality of images, the generating comprising, for each image:
8[b-i]	sending the image to a plurality of human observers, each human observer presented with a request to answer a question about a state of mind of a user in the image,
8[b-ii]	receiving, from each of the plurality of human observers, a response representing a judgment by the human observer of the state of mind of the user in the image,
8[b-iii]	generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers, and
8[b-iv]	storing the summary statistics in association with the image as part of the training data;
8[c]	training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image; and
8[d]	executing the trained model to predict a state of mind of a user in a new image.

Claim 15	Limitations
15[pre]	A computing system comprising: one or more processors; and a non-transitory computer readable storage medium, storing instructions that when executed by the one or more processors, cause the one or more processors to perform steps comprising:
15[a]	storing a plurality of images, each image displaying one or more users;
15[b]	generating training data from the plurality of images, the generating comprising, for each image:
15[b-i]	sending the image to a plurality of human observers, each human observer presented with a request to answer a question about a state of mind of a user in the image,
15[b-ii]	receiving, from each of the plurality of human observers, a response representing a judgment by the human observer of the state of mind of the user in the image,
15[b-iii]	generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers, and
15[b-iv]	storing the summary statistics in association with the image as part of the training data;
15[c]	training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image; and
15[d]	executing the trained model to predict a state of mind of a user in a new image.

III. LEVEL OF ORDINARY SKILL

Tesla alleges that a POSITA in the field of the '046 patent as of July 5, 2017 would have had “would have had at least: (1) a bachelor’s degree in electrical engineering, computer engineering, computer science, or equivalent course work, with three years of work experience in computer vision, autonomous vehicles, and/or machine learning; or (2) a master’s degree in electrical engineering, computer engineering, computer science, or equivalent course work, with a focus in computer vision, autonomous vehicles, and/or machine learning.” Pet. 12-13; *see also* Ex. 1003, ¶ 63. Neither Perceptive nor its expert (Dr. Saber) dispute this definition for purposes of resolving this IPR proceeding. Ex. 2039, ¶ 53.

Notably, Tesla does not allege that its expert (Dr. Janét) meets this definition, nor does Dr. Janét opine that he does. *See* Pet.; Ex. 1003. As such, Tesla has not established that Dr. Janét was a POSITA at the relevant time for purposes of this proceeding and thus his opinions should not be given consideration. *See, e.g., Kyocera Senco Indus. Tools Inc. v. Int’l Trade Comm’n*, 22 F.4th 1369, 1377-78 (Fed. Cir. 2022) (finding abuse of discretion when ALJ admitted “testimony on any issue that is analyzed through the lens of” a POSITA when expert did not qualify as POSITA). In contrast, Dr. Saber not only meets this definition today, but he also met it in the July 2017 timeframe. Ex. 2039, ¶ 53.

IV. CLAIM CONSTRUCTION

The term “user” recited in the challenged claims of the ’046 patent should be construed to mean “road user (e.g., motorist, bicyclist, pedestrian).”

The claims frequently recite images displaying a “user” and the state of mind of a “user.” *See, e.g.*, Ex. 1001 at claim 1 (“each image displaying one or more users,” “a state of mind of a user in the image,” “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image,” and “executing the trained model to predict a state of mind of a user in a new image”).

The specification makes clear that the claimed “user” in the context of images and a state of mind is a “road user (e.g., motorist, bicyclist, pedestrian).” *See, e.g., id.* at Abstract (“A computing device receives an image and a video segment of a road scene . . . response data is received, which includes at least one of an action and a likelihood of the action corresponding to another participant in the road scene.”), 3:58-60 (“Systems and methods are described for judging the behavior of people near a vehicle in a way that more closely resembles how human drivers would judge the behavior.”), 3:61-65 (“human observers[] view sample images of people (such as pedestrians) near streets and/or vehicles and indicate or are measured for their understanding of how they believe the people will behave”), 6:41-47 (“the statistics stored in the database 110 cover a large set of images of people on or near roads . . .

such as pedestrian, driver, motorcyclist, bicyclist, scooter driver, self-balancing scooter rider, unicyclist, motorized wheelchair user, skateboarder, or others”), 7:18-21 (“the prediction engine 114 uses the trained model from the model training system 112 to predict the actual, ‘real-world’ or ‘live data’ behavior of people on or near a road”), 7:44-50 (“The predictions produced by the trained model comprise a set of predictions of the state of mind of road users that can then be used to improve the performance of autonomous vehicles, robots, virtual agents, trucks, bicycles, or other systems that operate on roadways by allowing them to make judgments about the future behavior of road users based on their state of mind.”), 9:48-50 (“The judgment that the human observer makes is an evaluation of the state of mind of a road user depicted in the derived stimulus.”), 10:37-41 (“the pictured human road user has a given state of mind—e.g. that a pedestrian is likely to cross the street or that an oncoming vehicle is unlikely to be willing to yield to the vehicle containing the sensor if the vehicle containing the sensor needs to turn.”), 12:36-38 (“FIG. 9. is a flowchart showing a process of predicting the state of mind of road users using a trained learning algorithm.”), 13:50-54 (“The models described above can be implemented as a real-time module that makes predictions of road user behavior based on input from cameras or other sensors installed on a car 1000.”). There is no disclosure in the ’046 patent about the state of mind of a person that is not a road user. *See generally id.* Indeed, the title of the ’046 patent is “System and Method of

Predicting Human Interaction with Vehicles,” which refers to a predicting the state of mind of a road user. A POSITA reading the challenged claims in light of the specification would clearly understand the claimed “user” to mean a “road user (e.g., motorist, bicyclist, pedestrian).” Petitioner does not present any constructions for this term or for any other term. *See* Pet. at 13.

The Board need not construe any other claim terms in a particular manner to conclude that the challenged claims are not unpatentable. *See Wellman, Inc. v. Eastman Chem. Co.*, 642 F.3d 1355, 1361 (Fed. Cir. 2011) (claims “need only be construed ‘to the extent necessary to resolve the controversy’”).

V. APPLICABLE LEGAL STANDARDS

A. Tesla’s Burden

As the petitioner, Tesla bears the burden of showing the unpatentability of the challenged claims by a preponderance of the evidence. *See* 35 U.S.C. § 316(e). This burden never shifts to Perceptive. *See Harmonic Inc. v. Avid Tech., Inc.*, 815 F.3d 1356, 1363 (Fed. Cir. 2016) (“[T]he Petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable.”); *Amazon Web Services, Inc. v. Saint Regis Mohawk Tribe*, IPR2019-00103, Paper 22 at 17 (PTAB May 10, 2019) (“The burden is on Petitioner, not the Board, to specify with particularity how [the prior art] teaches [various claimed features].”); *Liberty Mut. Ins. Co. v. Progressive Cas. Ins. Co.*, No. CBM2012-00003, Paper 8 at 14-15

(P.T.A.B. Oct. 25, 2012) (representative) (“[I]t is the responsibility of the Petitioner to clearly articulate [the proposed combination].”); 37 C.F.R. §§ 42.22(a)(2), 42.104(b)(4). The Petition must provide an understandable explanation of the element-by-element specifics of the patentability challenges. *See* 35 U.S.C. § 312(a)(3); 37 C.F.R. §§ 42.22(a)(2), 42.104(b)(3)-(4).

B. Obviousness in IPR Proceedings

Each of Tesla’s grounds relies on obviousness. Section 103 of the Patent Act provides that “[a] patent may not be obtained . . . if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains.” 35 U.S.C. § 103(a). If a single element of the claim is absent from the prior art, the claim cannot be considered obvious. *See CFMT, Inc. v. YieldUp Int’l Corp.*, 349 F.3d 1333, 1342 (Fed. Cir. 2003) (“[O]bviousness requires a suggestion of all limitations in a claim.”) (citing *In re Royka*, 490 F.2d 981, 985 (C.C.P.A. 1974)); *In re Rijckaert*, 9 F.3d 1531, 1534 (Fed. Cir. 1993) (reversing obviousness rejection where prior art did not teach all claim limitations)).

Further, “[b]ecause an *inter partes* review may only be requested ‘on the basis of prior art consisting of patents or printed publications,’ . . . expert testimony cannot take the place of disclosure from patents or printed publications.” *See* OFFICE

PATENT TRIAL PRACTICE GUIDE at § I.G. “In other words, expert testimony may explain ‘patents and printed publications,’ but is not a substitute for disclosure in a prior art reference itself.” *Id.*; see *In re Kahn*, 441 F. 3d 977, 988 (Fed. Cir. 2006) (“[R]ejections on obviousness grounds cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness”).

C. The *Graham* Factors in IPR Proceedings

Obviousness is resolved on several factual determinations including: (1) the scope and content of the prior art, (2) any differences between the claimed subject matter and the prior art, and (3) the level of ordinary skill in the art. See *Graham v. John Deere Co.*, 383 U.S. 1, 17-18 (1966). These factors provide the threshold “background” for the obviousness determination, and without undertaking an analysis of the *Graham* factors, a proper obviousness determination cannot be made. *Id.* at 18 (“Under § 103, the scope and content of the prior art are to be determined; *differences between the prior art and the claims at issue are to be ascertained*; and the level of ordinary skill in the pertinent art resolved. *Against this background, the obviousness or nonobviousness of the subject matter is determined.*”) (emphasis added).

In view of this clear Supreme Court precedent, it is not surprising that IPR petitions “must address the *Graham* factors.” *Eizo Corp. v. Barco N.V.*, No.

IPR2014-00358, Paper 11 at 29-30 (P.T.A.B. July 23, 2014) (citing *Graham*, 383 U.S. at 17-18). For example, the Board faulted the petitioner in *Eizo* for failing to identify the differences between the claimed subject matter and the prior art. *Id.* at 29-30. In particular, the Board found insufficient the petitioner’s “conclusory assertion” that “[t]o the extent [the first reference] may not explicitly teach” the limitation, the second reference “explicitly teaches this limitation.” *Id.* at 30. The Board explained that “such an assertion fails to resolve the exact differences sought to be derived from” the second reference. *Id.* (finding that petitioner had not shown a reasonable likelihood of prevailing on that ground). Other Board decisions have reached the same result. *See, e.g., Moses Lake Indus., Inc. v. Enthone, Inc.*, No. IPR2014-00246, Paper 6 at 17 (P.T.A.B. June 18, 2014); *eBay, Inc. v. Paid, Inc.*, No. CBM2014-00125, Paper 15 at 21 (P.T.A.B. Sept. 30, 2014).

VI. THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 1

The Petition challenges claims 1-19 of the ’046 patent as obvious over Cox alone. *See* Pet. 13-34. However, as discussed in detail below, the challenged claims are not unpatentable in view of Cox alone not only because Cox fails to teach several

elements of the challenged claims, but also because Tesla failed to consider the *Graham* factors in the Petition.¹

A. Cox

Cox generally describes processes for improving machine learning systems for face detection in images. Ex. 1007, [0002] (describing “a machine-learning system trained to decide whether a face is present in an image will be given set of human-generated ‘ground truth’ labels that signal whether a human judged the image to contain a face or not”); *see also id.* (“In this case, the job of the machine-learning system is to emit analogous labels (for example, ‘face present’ and ‘face not present’) in response to new, previously unseen images.”). Cox uses this diametric “face present”/“face not present” example throughout its disclosure. *Id.* at [0002], [0008], [0010], [0012], [0013], [0016], [0017], [0026], [0029], [0033], [0039]-[0048], [0050], [0052]-[0056], Figs. 5, 6, 7, 9, 10, 11, 12 (all describing using Cox’s disclosure for face detection).

¹ In discussing the elements of the challenged claims that Cox does not disclose, Patent Owner references the Petition’s allegations for claim 1, because Tesla’s allegations for the corresponding elements of claims 8 and 15 merely refer back to its allegations for claim 1. *See* Pet. 31-34.

Cox explains that a limitation of machine learning at the time was “the unavailability of labeled data; it is difficult to provide an algorithm with enough labeled training data to achieve optimal performance.” *Id.*, [0003]. As Cox explains, “[w]ithout sufficiently large datasets, machine-learning algorithms tend to ‘overfit’ the data, adapting to spurious structure present in the training set that is not representative of the larger distribution of all examples in the real world.” *Id.*

To address this, Cox describes using “psychometric data” as an input in the machine learning system, in addition to the human annotations of the image. *See id.* at [0005]-[0010]. This psychometric data includes information about how long the human viewed the image before providing his or her response. *See id.* As Cox describes, this information can be used to determine the degree of “penalization” for misclassifying the image (i.e., indicating that the image contains a face when it does not, or vice versa), because “greater penalties may be introduced for misclassification . . . of examples that are easily classified by humans” and “lesser penalties may be introduced for misclassification of examples that are relatively more difficult for humans to classify.” *See id.*, [0005]. These penalties are used in a “human-weighted loss function” to train the model, with the aim of improving its accuracy. *See id.*, [0005]-[0010].

B. Cox Does Not Teach “user”/“users”

Elements [a], [b-i], [b-ii], [b-iii], [c], and [d] of each challenged independent claim recite a “user” or “users.” As discussed above, the Board should construe “user” to mean “road user (e.g., motorist, bicyclist, pedestrian).” *See* Section IV. Nothing in Cox teaches the claimed “user” under this construction. Ex. 2039, ¶¶ 71-73. Instead, the images described in Cox are of faces, and nowhere does Cox teach that those faces are those of a road user (e.g., motorist, bicyclist, pedestrian). *See, e.g.,* Ex. 1007, [0007]-[0010]; Ex. 2039, ¶ 73. Nowhere does the Petition allege that any of the images in Cox are of road users. *See* Pet. 13-34; Ex. 2039, ¶ 73. Cox therefore fails to teach the claimed “user”/“users.”

C. Cox Does Not Teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers”

Nothing in Cox teaches elements 1[b-iii], 8[b-iii], or 15[b-iii], which each require “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers.” The Petition alleges that several features of Cox meet the claimed “summary statistics,” but as explained below, each of those theories fails.

Psychometric Data: The Petition first alleges that Cox’s psychometric data constitutes “summary statistics.” Pet. 20-21. However, as Cox makes clear, its psychometric data “characterizes the annotation of the training objects by the

annotators” ([0018]) and that “may include or consist essentially of (i) response time for classifying one or more features, (ii) accuracy of feature classification, and/or (iii) presentation time of one or more training objects.” Ex. 1007, [0008]; *see also id.*, [0010], [0018], [0029]. None of those items (response time, accuracy, or presentation time) “describ[e] the state of mind of the user in the image” in any way. Ex. 2039, ¶ 77.

The Petition alleges that “Cox discloses generating ‘psychometric data’ . . . describing the emotions of a user in a sample training image.” Pet. 20. But that is inaccurate. Annotators’ response times and accuracy do not “describ[e] the emotions of a user in a sample training image” (*id.*), nor does the presentation time of the training object (i.e., the image) to the annotator. Ex. 2039, ¶ 77.

Tesla’s expert declaration adds nothing here, as its expert merely regurgitates the Petition’s allegations. *Compare* Pet. 20-21 *with* Ex. 1003, ¶ 77. *See Xerox Corp. v. Bytemark, Inc.*, IPR2022-00624, Paper 12 at 5 (P.T.A.B. Feb. 10, 2023) (affirming Board’s decision to “giv[e] little weight to Petitioner’s expert because the expert declaration merely offered conclusory assertions without underlying factual support and repeated, *verbatim*, Petitioner’s conclusory arguments”); *In re Ethicon, Inc.*, 844 F.3d 1344, 1348, 1352 (Fed. Cir. 2017) (concluding that the Board properly gave “little weight” to conclusory expert testimony of objective indicia); *TQ Delta, LLC v. Cisco Sys.*, 942 F.3d 1352, 1359 (Fed. Cir. 2019) (“This court’s opinions have

repeatedly recognized that conclusory expert testimony is inadequate to support an obviousness determination on substantial evidence review”).

Item-response curves: The Petition next alleges that Cox’s item-response curves constitute the claimed summary statistics. Pet. 22. Cox’s only description of “item-response curves” is that they are a “well-known technique[] . . . for measuring human performance,” that they measure “how consistent are judgment across a population,” and that “psychometric data may be displayed to analyze, e.g., response time per annotator.” Ex. 1007, [0006], [0029]; *see also* Fig. 3. As Cox’s Figure 3 shows, item-response curves plot the respondents’ ability and accuracy of their responses. *Id.*, Fig. 3; Ex. 2039, ¶ 84; Ex. 2041, Fig. 1. Item-response curves do not “describ[e] the state of mind of the user in the image.” Ex. 2039, ¶¶ 82-85.

Here too, Tesla’s expert declaration adds nothing, as Dr. Janét merely regurgitates the allegations in the Petition. *Compare* Pet. 22 *with* Ex. 1003, ¶ 78. *See Xerox*, Paper 12 at 5; *Ethicon*, 844 F.3d at 1348, 1352; *TQ Delta*, 942 F.3d at 1359.

Moreover, as discussed in Section VI.E below, Tesla’s failure to allege in the Petition that Cox’s item-response curves meet the “storing the summary statistics” requirement of elements 1[b-iv], 8[b-iv], and 15[b-iv] is fatal to any argument that Cox’s item-response curves can be the claimed “summary statistics.” As detailed

below, the “summary statistics” must be the same in the mapping of both claim elements, and Tesla does not map Cox’s item-response curves to both elements.

Human-weighted loss function: The Petition next alleges that Cox’s human-weighted loss function is the claimed summary statistics. Pet. 22-23. As Cox indicates, its human-weighted loss function is “computationally derived” “based at least in part on the classification data and the psychometric data.” Ex. 1007, [0007]; *see also id.*, [0008], [0009], [0019], [0036]. While the human-weighted loss function Cox describes can be derived based in part on the classification data, the human-weighted loss function itself does not “describ[e] the state of mind of the user in the image.” Instead, it is a computational result derived from the classification data and the psychometric data, and it does not itself describe the user in the image or his/her state of mind. Ex. 2039, ¶¶ 88-94. Tesla has not (and cannot) show that Cox’s human-weighted loss function satisfies the claimed summary statistics.

Dr. Janét’s declaration again regurgitates the allegations in the Petition, and his opinions thus do not alter this conclusion. *Compare* Pet. 22 *with* Ex. 1003, ¶ 79. *See Xerox*, Paper 12 at 5; *Ethicon*, 844 F.3d at 1348, 1352; *TQ Delta*, 942 F.3d at 1359.

Moreover, as discussed in Section VI.E below, Tesla’s failure to allege in the Petition that Cox’s human-weighted loss function meets the “storing the summary statistics” requirement of elements 1[b-iv], 8[b-iv], and 15[b-iv] is fatal to any

argument that Cox’s human-weighted loss function can be the claimed “summary statistics.” As detailed below, the “summary statistics” must be the same in the mapping of both claim elements, and Tesla does not map Cox’s human-weighted loss function to both elements.

End-of-test statistics: Finally, the Petition alleges that Cox’s end-of-test statistics can be the claimed summary statistics. Pet. 22-23. However, Cox’s only description of these end-of-test statistics is that they “may be displayed to the annotators 160 via display module 167.” Ex. 1007, [0048]; *see also* Fig. 8A. Nothing in Cox teaches that these end-of-test statistics “describ[e] the state of mind of the user in the image.” Cox does not describe the contents of these end-of-test statistics, an issue which alone is fatal to Tesla’s argument. Indeed, the mere statement that these statistics “may be displayed to the annotators,” and Figure 8A’s illegible content, conveys nothing about the content of those statistics. *See id.* at [0048], Fig. 8A. Moreover, because Cox teaches displaying these statistics to annotators, a POSITA would have expected them to contain information about the annotators’ performance and accuracy, not about the state of mind of any user in any image. Ex. 2039, ¶¶ 99-100.

Because Dr. Janét’s declaration again regurgitates the Petition’s allegations, his opinions do not impact this result. *Compare* Pet. 22-23 *with* Ex. 1003, ¶ 79. *See also Xerox*, Paper 12 at 5; *Ethicon*, 844 F.3d at 1348, 1352; *TQ Delta*, 942 F.3d at

1359. Moreover, at deposition, Dr. Janét confirmed that he could not discern the contents of Figure 8A and that Cox does not describe the content of these end-of-test statistics. Ex. 2045, pp. 160-164.

Additionally, as discussed in Section VI.E below, Tesla’s failure to allege in the Petition that Cox’s end-of-test statistics meet the “storing the summary statistics” requirement of elements 1[b-iv], 8[b-iv], and 15[b-iv] is fatal to any argument that Cox’s end-of-test statistics can be the claimed “summary statistics.” As detailed below, the “summary statistics” must be the same in the mapping of both claim elements, and Tesla does not map Cox’s end-of-test statistics to both elements.

Each of Tesla’s theories therefore fail, and Cox does not teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers,” as required by elements 1[b-iii], 8[b-iii], and 15[b-iii].

D. Cox Does Not Teach “storing the summary statistics in association with the image as part of the training data”

Tesla has not (and cannot) show that Cox teaches elements 1[b-iv], 8[b-iv], and 15[b-iv], which each require “storing the summary statistics in association with the image as part of the training data.”

Critically, “*the summary statistics*” elements 1[b-iv], 8[b-iv], and 15[b-iv] derive antecedent basis from “generating *summary statistics*” of elements 1[b-iii],

8[b-iii], and 15[b-iii]. Tesla fails to appreciate or acknowledge this antecedent-basis linkage, which leads to problems.

As discussed above, for elements 1[b-iii], 8[b-iii], and 15[b-iii], Petitioner argues that summary statistics is met by Cox's (i) psychometric data, (ii) human-weighted loss function, or (iii) end-of-test statistics. For elements 1[b-iv], 8[b-iv], and 15[b-iv], however, the Petition only points to storing Cox's psychometric data as meeting the "summary statistics" element. Pet. 23-24. There is no mention of either Cox' human-weighted loss function or Cox's end-of-test statistics allegedly meeting the "summary statistics" of elements 1[b-iv], 8[b-iv], and 15[b-iv]. *See id.*; Ex. 2039, ¶ 109. And even if Tesla had, Cox does not teach storing any of those items "in association with the image as part of the training data," as required by the challenged claims. Ex. 2039, ¶ 110.

As discussed in Section VI.C above, Cox's psychometric data does not satisfy the claimed summary statistics of elements 1[b-iii], 8[b-iii], and 15[b-iii] because it fails to meet the requirement that "summary statistics describ[e] the state of mind of the user in the image." Because of the antecedent-basis linkage, the fact that Cox's psychometric data does not meet elements 1[b-iii], 8[b-iii], and 15[b-iii] is likewise

fatal to any argument that such psychometric data can meet the requirements of elements 1[b-iv], 8[b-iv], and 15[b-iv].

Relatedly, because Petitioner only points to Cox's psychometric data as meeting the summary statistics of elements 1[b-iv], 8[b-iv], and 15[b-iv], the Board need not consider whether Cox's human-weighted loss function or end-of-test statistics meet the "summary statistics" of elements of 1[b-iii], 8[b-iii], and 15[b-iii]. The claimed summary statistics are the same in both sets of claim elements and, in absence of an obviousness theory that mixes and matches Cox's embodiments that Petitioner has not even attempted to articulate, Petitioner must point to the same information from Cox (i.e., either psychometric data, human-weighted loss function, or end-of-test statistics) as satisfying this requirement. Therefore, not only does Cox not teach elements 1[b-iv], 8[b-iv], or 15[b-iv], Tesla's lone theory that psychometric data meets the claimed summary statistics of these elements dooms its other theories for elements 1[b-iii], 8[b-iii], and 15[b-iii].

E. Cox Does Not Teach "training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image"

Cox does not teach elements 1[c], 8[c], or 15[c], which each require "training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image."

The only thing the Petition alleges that Cox’s model predicts is “classifications for the training images using the model.” Pet. 25. However, classifications in Cox do not satisfy the claimed “summary statistics.” Instead, Cox’s classifications are merely the model’s prediction of how to classify the image (i.e., whether it contains a face or not). *See* Ex. 1007 at, e.g., [0033], Fig. 5; Ex. 2039, ¶ 115. This flaw is fatal to Ground 1. Moreover, the Petition and Dr. Janét’s declaration’s conclusory allegations do not explain how or why Cox’s classifications satisfy the summary statistics element. *See id.*; *see* Ex. 1003, ¶¶ 81-83. Accordingly, Tesla has not met its burden. *See Harmonic*, 815 F.3d at 1363; *Amazon*, IPR2019-00103, Paper 22 at 17; *Liberty Mut.*, No. CBM2012-00003, Paper 8 at 14-15.

Moreover, while the Petition alleges elsewhere that other features of Cox constitute “summary statistics,” the Petition does not do so for elements 1[c], 8[c], or 15[c]. *Compare* Pet. 20-23 *with* Pet. at 24-27. Tesla has therefore waived any such arguments and must rely solely on its flawed argument that classifications in Cox satisfy the claimed summary statistics of these elements. *See, e.g., In re NuVasive, Inc.*, 841 F.3d 966, 971-72 (Fed. Cir. 2016) (vacating Board’s unpatentability finding based on portion of reference not raised in petition); *see also, e.g., Juniper Networks, Inc. v. Swarm Tech. LLC*, IPR2021-01445, Paper 62 at 28-29 (P.T.A.B. Feb. 7, 2023) (refusing to consider arguments raised in petitioner’s reply but not contained in petition). However, even if Tesla had not waived those arguments, those

additional features of Cox (i.e., its psychometric data, item-response curves, human-weighted loss function, or end-of-test statistics) do not satisfy the claimed summary statistics, as nothing in Cox describes its model being configured to predict any of those items. Ex. 2039, ¶ 116.

In addition, Cox does not “train[] a model using the training data, the model configured to . . . predict summary statistics.” Instead, Cox trains a “classifier” using a human-weighted loss function to classify features (e.g., “face present” or “face not present”). The output of Cox’s model is a binary or categorical classification, not a set of summary statistics describing a person’s state of mind. The Petition’s argument that Cox’s model optimization equates to predicting summary statistics (*see* Pet. 24-25) is a fundamental mischaracterization of what Cox’s model outputs. *See* Ex. 1007, [0035]-[0036] (describing using the concept of regularization to yield an alternative model with a human-weighted loss function (formula at bottom of [0036]) and using that instead of the formula at bottom of [0035] for training).

For these reasons, Cox does not teach “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image,” as required by elements 1[c], 8[c], and 15[c].

F. Challenged Dependent Claims

As discussed above, Cox does not teach several elements of claims 1, 8, and 15. Each of the remaining challenged claims in Ground 1 depend from one of those independent claims, and Cox does not render them obvious for the same reasons.

G. Tesla and its Expert Fail to Analyze the *Graham* Factors, As Required for a Proper Obviousness Analysis

Even aside from the Petition’s failures (discussed above) to demonstrate how Cox teaches several elements of the challenged claims, the Petition commits a fundamental legal error in Ground 1, as it fails to consider the *Graham* factors, most notably the requirement of considering “any differences between the claimed subject matter and the prior art.” *Graham*, 383 U.S. at 18.

Because Ground 1 is a single-reference obviousness assertion (as opposed to an anticipation theory), Petitioner presumably proposes to modify Cox to arrive at the challenged claims. *See, e.g., Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1361 (Fed. Cir. 2016) (“Though less common, in appropriate circumstances, a patent can be obvious in light of a single prior art reference if it would have been obvious to modify that reference to arrive at the patented invention.”). Both Tesla and its expert fail to articulate differences between the challenged claims and Cox, and neither specifies a manner in which they propose to modify Cox to arrive at the challenged claims. *See generally* Pet. 13-34; Ex. 1003, ¶¶ 66-120. Without identifying these

differences, a proper obviousness determination cannot be made. *Graham*, 383 U.S. at 17-18; *Eizo*, Paper 11 at 29-30; *Moses Lake*, Paper 6 at 17; *eBay*, Paper 15 at 21.

At his deposition, Dr. Janét could not explain any differences between Cox and the challenged claims, other than to point to statements in his declaration where he believed a person who was not a POSITA could misunderstand Cox's disclosure. Ex. 2045, pp. 98-115. Such allegations are not only irrelevant to the obviousness determination (which is performed from the perspective of a POSITA), they are not contained in the Petition or Dr. Janét's declaration and have thus been waived. 35 U.S.C. § 103 (obviousness inquiry concerns "the differences between the claimed invention and the prior art . . . to a person having ordinary skill in the art to which the claimed invention pertains"); *see also, e.g., NuVasive*, 841 F.3d at 971-72; *Juniper*, Paper 62 at 28-29.

Given this legal failure of the Petition, the Board should find the challenged claims not unpatentable under Ground 1 for this additional reason.

VII. THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 2

Ground 2 of the Petition challenges claims 1-19 of the '046 patent as obvious over Cox in view of Ross. *See generally* Pet. 34-45. However, as discussed below, the challenged claims are not unpatentable over Cox in view of Ross not only because Cox fails to teach several elements of the challenged claims, but also

because of a lack of motivation to combine Cox and Ross in the manner Tesla alleges.²

A. Ross

In contrast to Cox’s focus on face detection (*see* Section VI.A), Ross concerns predicting what an object in a road scene will do. Ex. 2039, ¶ 124 (citing Ex. 1005, 1:30-54, Fig. 10). Also, while Cox uses crowd-sourced human annotation data to train its model, Ross’s training data is based on actual road user behavior over time. Ex. 2039, ¶ 124 (citing Ex. 1005, 1:30-54, Fig. 10). These fundamental differences between Cox and Ross demonstrate why a POSITA would not have been motivated to combine Cox and Ross in the manner Tesla alleges. Further, Tesla’s use of Ross is limited, and it does not cite to Ross to cure the shortcomings of Cox that were discussed in Section VI for Ground 1.

B. Cox and Ross Do Not Teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers”

As discussed above, Cox does not teach elements 1[b-iii], 8[b-iii], or 15[b-

² In discussing the claim elements herein that the art does not teach, Patent Owner references the Petition’s allegations for claim 1, because Tesla’s allegations for the corresponding elements of claims 8 and 15 throughout the Petition merely refer back to allegations for claim 1. *See, e.g.*, Pet. 45.

iii], which each require “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers.” For Ground 2, the Petition merely refers back to its Ground 1 (*i.e.*, Cox alone) arguments, and it does not allege that Ross teaches these elements. Pet. 42. Thus, the combination of Cox and Ross fails to teach these elements for the same reasons as discussed above with respect to Ground 1.

C. Cox and Ross Do Not Teach “storing the summary statistics in association with the image as part of the training data”

As discussed above for Ground 1, Cox does not teach elements 1[b-iv], 8[b-iv], or 15[b-iv], which each require “storing the summary statistics in association with the image as part of the training data.” For Ground 2, the Petition merely refers to Ground 1 (*i.e.*, Cox alone), and it does not allege that Ross teaches these elements. Pet. 42. Thus, Cox and Ross fail to teach these elements for the same reasons as discussed above with respect to Ground 1.

D. Cox and Ross Do Not Teach “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image”

Cox and Ross do not teach elements 1[c], 8[c], or 15[c], which each require “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image.” With respect to this element, the Petition merely alleges

that “[a]s combined, the training of Cox’s model would have been performed as discussed in Section X.A.2.h (element 1[c] – Ground 1) except the training data includes training images and annotations with respect to actions by a person in images reflecting possible intents of the person, and Cox’s model is trained on such training data as discussed in Section X.A.2.h (element 1[c] – Ground 1).” Pet. 42-43. Nothing in the cross-referenced section of the Petition (which relates to Ground 1) alleges that anything in Ross meets this element. *See* Pet. 24-27. Instead, there the Petition alleges that Cox predicts “classifications for the training images using the model,” which as discussed above are not “summary statistics.” The Petition’s allegation that “[a]s combined, . . . the training data includes training images and annotations with respect to actions by a person in images reflecting possible intents of the person” (Pet. 42-43) likewise does not allege anything that the models of Cox or Ross predict.

Cox and Ross therefore do not teach “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image,” as required by elements 1[c], 8[c], and 15[c]. Tesla has failed to carry its burden on this point.

E. Challenged Dependent Claims

As discussed above, Cox and Ross do not teach several elements of claims 1, 8, and 15. Each of the remaining challenged claims in Ground 2 depend from one of

those independent claims, and Cox and Ross do not render them obvious for the same reasons.

F. A POSITA Would Not Have Been Motivated to Combine Cox and Ross in the Manner Tesla Alleges

As best the Petition can be understood, Tesla proposes to combine Cox with Ross for elements 1[b-i], 1[b-ii], and 1[c]. Pet. 34-35. Petitioner's lone modification of Cox based on Ross appears to be that, in view of Ross, the "user" in the image of Cox can be "a person or road user." Pet. 36. Nowhere does the Petition describe how a POSITA would have accomplished this modification of Cox in view of Ross, or why a POSITA would have been motivated to make such proposed modification.

As discussed above, Cox is generally directed to determining the presence or absence of a face in images, whereas Ross is directed to predicting what an object in a road scene will do. As Dr. Saber explains, Cox's machine learning system is specifically designed for the tasks it is intended to accomplish, including the design of the network architecture, the regularization methodology applied, and the selection of the particular loss function (a human weighted loss function instead of a hinge loss function). Ex. 2039, ¶¶ 149-152. Given that Ross's system is designed for a completely different purpose and is instead intended to predict intents and behaviors of road users (and is specifically designed to do so by looking into the future in a dataset to see what the actual result was), Ross's system requires a different system design with different training and validation processes and different datasets. *Id.*

As to Tesla’s first alleged motivation to combine (i.e., “transmit[ting] a training image to multiple human annotators for annotation with a prompt to answer about an action by a person or road user in the image reflecting a possible intent of the person”), Ross is not directed to transmitting images to crowdsourced human annotators. Thus, a POSITA would not have been motivated to combine Cox with Ross to arrive at this modification. Ex. 2039, ¶ 153.

As to Tesla’s second alleged motivation to combine (i.e., training Cox’s system “to analyze an action by a person reflecting a possible intent in images, as taught by Ross”), here too, a POSITA would not have been motivated to combine Cox with Ross. Cox is generally concerned with predicting classifications regarding human faces, whereas Ross is generally concerned with predicting what road users will do by observing behavior past the relevant point in time. Ex. 2039, ¶ 154 (citing Ex. 1007, [0016], [0048], [0049], Ex. 1005, 4:25-32, 4:51-65). Given these fundamental differences, a POSITA would not have been motivated to modify Cox’s human-annotation driven face-recognition/classification system to incorporate Ross’s technologically distinct mechanism of using road user behavior over time to train a model to predict road user behavior. Ex. 2039, ¶ 154.

Neither the Petition nor Dr. Janét account for the fundamental differences between Cox and Ross. Both the Petition and Dr. Janét simply allege, in conclusory fashion and with no supporting analysis or citations, that their combination is proper.

Tesla and Dr. Janét’s assertion that “Cox discloses a machine learning prediction system that relies on human annotators to annotate training images regarding emotional states of a person in the images” and that “Ross provides a well-known application for such a system by teaching that the subjects in the images are road users . . . performing an action on the road” (Pet. 37; Ex. 1003, ¶ 127) ignores the technical and practical differences between these systems. Tesla and Dr. Cox’s surface-level allegations that both Cox and Ross are intended to “make human-like prediction” (Pet. 37-38; Ex 1003, ¶¶ 128-129) likewise ignores those differences. Nor do the conclusory assertions that “combining Ross with Cox’s teachings about generating and using the same data would have required minimal changes,” that “[m]aking this combination would also be well within the skill of a POSITA,” or that “a POSITA would have had a reasonable expectation of success” (Pet. 39-40; Ex. 1003, ¶ 130) account for those differences.

At bottom, Tesla and Dr. Janét’s motivation to combine amounts to an assertion that both references relate to “human-like” predictions, but this generalization is far too broad. Nearly every machine learning reference in the autonomous vehicle space aspires to “human-like” performance. Ex. 2039, ¶ 150; Ex. 2044 at p. VII (“The grand vision of neural networks is to create artificial intelligence by building machines whose architecture simulates the computations in the human nervous system.”). The specific combination—using Cox’s crowd-

sourced psychometric annotation framework with Ross’s autonomous vehicle intent prediction—requires hindsight reconstruction of the claimed invention. *TQ Delta*, 942 F.3d at 1361 (warning against “us[ing] the challenged patent as a roadmap to reconstruct the claimed invention using disparate elements from the prior art”).

Moreover, a POSITA would have been aware of the unique nature of the problems faced by attempting to analyze image, video, or other sensor data in real time driving situations. A POSITA would not have been motivated combine proposed solutions for driving situations with Cox’s crowdsourced annotation system for improving face detection, which does not present the same unique challenges. Ex. 2039, ¶ 156; Ex. 2042 at 1 (noting “pedestrian path prediction is a challenging problem, due to the highly dynamic nature of pedestrian motion, and systems need to react with limited computation time”); Ex. 2043 at 1 (noting that “the ability to infer the future intentions of drivers to predict the likelihood of potential collisions” “is a challenging task, particularly in intersections where complex traffic scenarios result in a proportionally high number of accidents”).

For these reasons, Tesla has failed to show that a POSITA would have been motivated to combine Cox and Ross in the manner Tesla and Dr. Janét allege, and the challenged claims should be found not unpatentable under Ground 2 for this reason.

VIII. THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 3

Ground 3 challenges dependent claims 5, 12, and 19 as being obvious over Cox in view of Ellenbogen. *See* Pet. 46-49. For this ground, the Petition does not allege that Ellenbogen teaches any elements of the independent claims. *See id.* As discussed above with respect to Ground 1, Cox does not teach several elements of the independent claims. Ground 3 therefore fails to render obvious any of the challenged claims for the same reasons as discussed above for Ground 1.

IX. THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 4

Ground 4 of the Petition challenges claims 1-19 of the '046 patent as obvious over Cox in view of Munro. *See generally* Pet. 49-54. However, as discussed below, the challenged claims are not unpatentable over Cox in view of Munro not only because Cox and Munro fail to teach several elements of the challenged claims, but also because of a lack of motivation to combine Cox and Munro in the manner Tesla alleges.³

³ In discussing the elements of the challenged claims that Cox and Munro do not disclose, Patent Owner references the Petition's allegations for claim 1, because Tesla's allegations for the corresponding elements of claims 8 and 15 merely refer back to its allegations for claim 1. *See* Pet. 54.

A. Munro

In contrast to Cox and the '046 patent, Munro is generally directed to natural language processing. Thus, instead of being concerned with analyzing the contents of an image for a given purpose (i.e., a face present/face not present determination as described in Cox, or determining the intent of a road user in an image as described in the '046 patent), Munro is directed to improving processes for analyzing the content of documents. *Compare, e.g.,* Ex. 1007, [0007]-[0010] *and* Ex. 1001, 1:27-61 *with, e.g.,* Ex. 1006, Abstract, [0003], [0004]; *see also* Ex. 2039, ¶¶ 165-167.

B. Cox and Munro Do Not Teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers”

Nothing in Cox or Munro teaches elements 1[b-iii], 8[b-iii], or 15[b-iii], which each require “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers.”

The Petition alleges that “[t]he combination of Cox and Munro teaches that Cox’s generated statistical data from responses of human annotators regarding the emotional state of a user in a training image would have included tallies to aggregate the number of annotators who similarly classified the emotional state of the user in

the training image.”⁴ Pet. 52-53. However, the “tallies” the Petition refers to from Munro are part of an “aggregation process” that is used to “analyze[] all of the annotations made to a single document to determine whether said document may be reliably annotated for use in training the natural language model.” Ex. 1006, [0140]. Nothing in Munro suggests that its tallies “describ[e] the state of mind of the user in the image.” Indeed, nothing in Munro suggests that its tallies have anything to do with an image, because Munro’s disclosure concerns analysis of documents, not images. *See* Ex. 1006, [0140]; Ex. 2039, ¶¶ 172-179. In addition, as Munro expressly states, this aggregation process is intended to determine whether a given document should remain in the training data—a process wholly different than “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers.” *See* Ex. 1006, [0140]; Ex. 2039, ¶¶ 172-179.

The combination of Cox and Munro therefore fails to teach “generating summary statistics describing the state of mind of the user in the image based on the

⁴ As set forth in Ground 1, Cox does alone does not disclose “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers.”

received responses from the plurality of human observers,” as required by elements 1[b-iii], 8[b-iii], and 15[b-iii].

C. Cox and Munro Do Not Teach “storing the summary statistics in association with the image as part of the training data”

Nothing in Cox or Munro teaches elements 1[b-iv], 8[b-iv], or 15[b-iv], which each require “storing the summary statistics in association with the image as part of the training data.”

With respect to this element, the Petition alleges that “Cox and Munro teach [it] for at least the same reasons discussed in Section X.A.2.g (element 1[b-iv]), except the generated summary statistics as taught by Munro would have been stored in Cox’s computer record in a database for the statistic’s associated image.” Pet. 53-54. This theory fails for two independent reasons.

First, the Petition does not expressly indicate what it means by “the generated summary statistics as taught by Munro,” leaving Patent Owner and the Board to guess as to what aspect of Munro constitutes “generated summary statistics.” *See Harmonic*, 815 F.3d at 1363 (“[T]he Petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable.”); *Liberty Mut.*, Paper 8 at 14-15 (PTAB Oct. 25, 2012); *see also* 37 C.F.R. §§ 42.22(a)(2), 42.104(b)(4). The only plausible reading in view of Tesla’s allegations for element 1[b-iii] is that this phrase refers to “tallies to aggregate the number of annotators who similarly classified the emotional state of the user in the training image.” Pet. 53. However,

as discussed above, the tallies described in Munro are part of an “aggregation process” that is used to “analyze[] all of the annotations made to a single document to determine whether said document may be reliably annotated for use in training the natural language model.” Ex. 1006, [0140]. Thus, not only does nothing in Munro suggest that its tallies are stored anywhere, Munro indicates that its tallies are intended to identify documents to *exclude* from the training data.

Second, the Petition likewise does not indicate what is meant by its allegation that “the generated summary statistics as taught by Munro would have been stored in Cox’s computer record in a database for the statistic’s associated image,” including what “Cox’s computer record” is or what “a database for the statistic’s associated image” is. *See Harmonic*, 815 F.3d at 1363; *Liberty Mut.*, Paper 8 at 14-15 (PTAB Oct. 25, 2012); *see also* 37 C.F.R. §§ 42.22(a)(2), 42.104(b)(4). The only plausible reading in view of the Petition’s cross-referenced allegations in Section X.A.2.g of the Petition is that Tesla proposes to store Munro’s tallies in Cox’s “database of training objects.” Pet. 23. Nothing in Munro teaches storing these tallies in any database, whether with the training data or otherwise. Also, as discussed above, the purpose of Munro’s aggregation process is to determine whether to exclude a document from the training data. Ex. 1006, [0140]. Thus, for documents that the tallies indicate “that this document[] is too ambiguous to be used in the adaptive machine learning training process,” Munro suggests *excluding* such

documents from the training dataset—which would suggest removing or deleting the document (and any associated tallies) from the dataset. *Id.*; Ex. 2039, ¶ 186.

Therefore, the combination of Cox and Munro does not teach “storing the summary statistics in association with the image as part of the training data,” as required by elements 1[b-iv], 8[b-iv], and 15[b-iv]. Additionally, Tesla’s vague allegations fall woefully short of meeting its burden.

D. Cox and Munro Do Not Teach “the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image”

Nothing in Cox or Munro teaches elements 1[c], 8[c], or 15[c], which each require that “the model [be] configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image.”

The Petition alleges that, “[a]s combined, the training of Cox’s model would have been performed as discussed in Section X.A.2.h (element 1[c]) except the training data includes the generated summary statistics as taught by Munro, and Cox’s model is trained on such training data in the same manner discussed in Section X.A.2.h (element 1[c]).” Pet. 54.

As with elements 1[b-iv], 8[b-iv], and 15[b-iv], the Petition does not expressly indicate what it means by “the generated summary statistics as taught by Munro,” leaving Patent Owner and the Board to guess as to what constitutes “generated summary statistics.” *See Harmonic*, 815 F.3d at 1363; *Liberty Mut.*, Paper 8 at 14-

15 (PTAB Oct. 25, 2012); *see also* 37 C.F.R. §§ 42.22(a)(2), 42.104(b)(4). But as stated above, the only plausible reading is that this refers to “tallies to aggregate the number of annotators who similarly classified the emotional state of the user in the training image.” Pet. 53. And as discussed above, Munro’s tallies are part of an “aggregation process” that is used to “analyze[] all of the annotations made to a single document to determine whether said document may be reliably annotated for use in training the natural language model.” Ex. 1006, [0140]. Nothing in either Cox or Munro suggests “predict[ing]” any such tallies, as expressly required by element [c]. Ex. 2039, ¶ 194.

Therefore, the combination of Cox and Munro does not teach “the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image,” as required by elements 1[c], 8[c], and 15[c]. Additionally, Tesla’s vague allegations fall woefully short of meeting its burden.

E. Challenged Dependent Claims

As discussed above, Cox and Munro do not teach several elements of claims 1, 8, and 15. Each of the remaining challenged claims in Ground 4 depend from one of those independent claims, and Cox and Munro do not render them obvious for the same reasons.

F. A POSITA Would Not Have Been Motivated to Combine Cox and Munro in the Manner Tesla Alleges

Even setting aside the fact that Cox and Munro fail to teach several of the challenged claims' elements, Ground 4 suffers from an additional flaw in that a POSITA would not have been motivated to combine Cox and Munro in the manner Tesla and Dr. Janét allege.

As discussed above, Cox and Munro relate to different fields—Cox to object recognition in images, and Munro to natural language processing of documents. *Compare, e.g.,* Ex. 1007, [0007]-[0010] *with, e.g.,* Ex. 1006, Abstract, [0003], [0004]; *see also* Ex. 2039, ¶¶ 165-167. As Dr. Saber explains, the differences between the fields of Cox and Munro go far beyond just the types of inputs they consider and include fundamental differences in what they aim to accomplish, the technological processes used to process those inputs, and the form of their outputs. Ex. 2039, ¶¶ 165-167, 172-179. Nothing in the Petition or Dr. Janét's declaration acknowledges much less accounts for these fundamental differences. *See* Pet. 49-54; Ex. 1003, ¶¶ 160-176. And, these differences relate directly to Tesla's proposal to combine Cox and Munro, because Munro's aggregation process (which is the crux of Tesla's proposed combination) relates specifically to an issue in natural language processing, i.e., whether a "document may be reliably annotated for use in training the natural language model." Ex. 1006, [0140]. As Dr. Saber explains, performing such an aggregation process and excluding data in an autonomous driving context

would not be practical. Ex. 2039, ¶¶ 172-179. Thus, the differing practical realities of Cox’s and Munro’s systems and applications necessarily teaches away from combining them in the manner Tesla suggests. Ex. 2039, ¶¶ 172-179, 198.

The lack of any motivation to combine Cox and Munro is an independent reason that the challenged claims are not unpatentable under Ground 4.

G. The Petition Fails to Analyze the *Graham* Factors, As Required for a Proper Obviousness Analysis

Finally, the Petition fails to consider the *Graham* factors in its obviousness analysis. In fact, in discussing Ground 4, the Petition is contradictory, because in the same paragraph the Petition alleges that “Cox alone renders obvious the Challenged Claims,” it also alleges that “Cox is further combined with Munro regarding the following limitation in Ground 4: generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers.” Pet. 49. This express failure to account for the differences between the challenged claims and the art is not only precisely what *Graham* requires but it is also what the Board has repeatedly rejected in past decisions. *Graham*, 383 U.S. at 17-18; *Eizo*, Paper 11 at 29-30; *Moses Lake*, Paper 6 at 17; *eBay*, Paper 15 at 21.

X. THE CHALLENGED CLAIMS ARE NOT UNPATENTABLE UNDER GROUND 5

Ground 5 of the Petition challenges claims 1-19 of the '046 patent as obvious over Ellenbogen in view of Munro. *See generally* Pet. 55-76. However, as discussed in detail below, the challenged claims are not unpatentable over Ellenbogen in view of Munro not only because Ellenbogen and Munro fail to teach several elements of the challenged claims, but also because of a lack of motivation to combine Ellenbogen and Munro in the manner Tesla alleges.⁵

A. Ellenbogen

Ellenbogen generally describes the design of a “[p]latform for gathering real-time analysis.” Ex. 1004, Title. Ellenbogen describes “a predictive model of a machine computation component [that] can be trained by providing the sensor data as input and the result from the agent computation component as a supervisory signal.” Ex. 1004, [0136]. The portions of Ellenbogen the Petition focuses on describe the use of a convolutional neural network to make “binary or multiclass” classifications of objects in images. *See id.*, [0163]-[0166].

⁵ In discussing the elements of the challenged claims that Ellenbogen and Munro do not disclose, Patent Owner references the Petition’s allegations for claim 1, because Tesla’s allegations for the corresponding elements of claims 8 and 15 merely refer back to its allegations for claim 1. *See* Pet. 73-76.

In contrast, as discussed above, Munro is generally directed to natural language processing, and specifically to improving processes for analyzing the content of documents. *See* Section IX.A.

B. Ellenbogen and Munro Do Not Teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers”

Nothing in Ellenbogen or Munro teaches elements 1[b-iii], 8[b-iii], or 15[b-iii], which each require “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers.”

The Petition alleges that “Munro teaches that Ellenbogen’s composite output of query result responses regarding the behavior or activity of a person in a sample image would have included an aggregate of the number of annotators who classified the behavior or activity as having various labels in their responses.” Pet. 64. For starters, Tesla’s characterization of Ellenbogen is inaccurate. Ellenbogen’s composite output is merely a plurality of results that are derived from requesting processing by more than one agent. Ex. 1004, [0135], [0149]. Ellenbogen does not teach an “aggregate” of any kind, and the Petition does not allege that Ellenbogen teaches anything that meets the claimed requirement of “generating summary statistics describing the state of mind of the user in the image.” *See id.* at 64-65. As best the Petition can be understood, Tesla proposes to combine Ellenbogen with

Munro such that Ellenbogen’s composite output includes the tallies that Munro teaches, which the Petition calls the “Ellenbogen-Munro composite output.” *See* Pet. 55-56.

However, Munro’s tallies are not “summary statistics describing the state of mind of the user in the image.” Munro expressly indicates that the purpose of this aggregation is to “analyze[] all of the annotations made to a single document to determine whether said document may be reliably annotated for use in training the natural language model.” Ex. 1006, [0140]. Nothing in Munro suggests that its tallies “describe[e] the state of mind of the user in the image.” Nothing in Munro suggests that its tallies have anything to do with an image, because Munro’s disclosure concerns analysis of documents, not images. *See* Ex. 1006, [0140]; Ex. 2039, ¶ 204. Indeed, Munro expressly states that its aggregation process is intended to determine whether a given document should remain in the training data—a process wholly different than “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers.” *See* Ex. 1006, [0140]; Ex. 2039, ¶ 204.

The combination of Ellenbogen and Munro therefore fails to teach “generating summary statistics describing the state of mind of the user in the image based on the received responses from the plurality of human observers,” as required by elements 1[b-iii], 8[b-iii], and 15[b-iii].

C. Ellenbogen and Munro Do Not Teach “storing the summary statistics in association with the image as part of the training data”

Nothing in Ellenbogen or Munro teaches elements 1[b-iv], 8[b-iv], or 15[b-iv], which each require “storing the summary statistics in association with the image as part of the training data.”

With respect to this element, the Petition alleges that “Munro teaches to store the Ellenbogen-Munro composite output in association with the image as part of the training data” because “Munro’s system includes a ‘client data store’ (i.e., a database of training data) that contains ‘a repository of data that is used to train and ultimately generate the natural language model.’” Pet. 66. But, nothing in any of these disclosures teaches storing Munro’s tallies from its aggregation process in any location, let alone “in association with the image as part of the training data.” Ex. 2039, ¶ 211. Instead, as discussed above, the tallies described in Munro are part of an “aggregation process” that is used to “analyze[] all of the annotations made to a single document to determine whether said document may be reliably annotated for use in training the natural language model.” Ex. 1006, [0140]. Thus, not only does Munro not teach or suggest storing these tallies in any particular location, but Munro expressly indicates that its tallies are intended to identify documents to *exclude* from the training data. *Id.*, Ex. 2039, ¶ 211.

Therefore, the combination of Cox and Munro does not teach “storing the summary statistics in association with the image as part of the training data,” as required by elements 1[b-iv], 8[b-iv], and 15[b-iv].

D. Ellenbogen and Munro Do Not Teach “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image”

Nothing in Ellenbogen or Munro teaches elements 1[c], 8[c], or 15[c], which each require “training a model using the training data, the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image.”

As discussed above, the so-called “Ellenbogen-Munro composite output” is not “summary statistics describing the state of mind of the user.” *See* Section X.B. Perhaps because of this shortcoming, here Tesla switches its theory for what constitutes the claimed “summary statistics describing the state of mind of the user” versus what it alleged in elements 1[b-iii]/[b-iv], 8[b-iii]/[b-iv], and 15[b-iii]/[b-iv]. Specifically, instead of alleging that the “Ellenbogen-Munro composite output” is the claimed “summary statistics describing the state of mind of the user,” here Tesla instead alleges that “predicting a multiclass behavior or activity prediction of a person in a training image and associated confidence scores” teaches the claimed “predict[ing] summary statistics describing the state of mind of the user.” *Compare* Pet. 64 *with id.* at 69.

First, a POSITA would understand the plain and ordinary meaning of “summary statistics” in elements 1[c], 8[c], and 15[c] to refer to the same summary statistics in elements 1[b-iii]/[b-iv], 8[b-iii]/[b-iv], and 15[b-iii]/[b-iv]. Ex. 2039, ¶ 217. This is true because those elements relate to “generating training data,” and elements 1[c], 8[c], and 15[c] relate to “training a model using the training data.” As Dr. Saber explains, a POSITA would not interpret the challenged claims in the manner Tesla alleges—which would be to recite generating one set of summary statistics in the “[b]” elements but having a model be configured to predict a different set of summary statistics. *Id.* Thus, Tesla’s effort to read the “summary statistics” here onto a different feature of Ellenbogen defies the plain and ordinary meaning of the language of the challenged claims and should be rejected. *Id.*

Second, even setting aside Tesla’s improper interpretation of the claims, the “multiclass behavior or activity prediction” Tesla identifies does not meet the claims’ requirement that “the model [be] configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image.” The portion of Ellenbogen that Tesla and Dr. Janét rely on merely indicates that “[t]he object classifier can generate a score from 0.0 to 1.0 related to the confidence that the object is not present (0.0) or present (1.0)” and that “[t]he classification can be binary or multiclass.” *See* Pet. 68 (citing Ex. 1004, [0166]). Generating such object present/not present classifications (even with

associated confidence values) is not “predict[ing] summary statistics describing a state of mind of the user in the input image,” as those object present/not present classifications do not “describ[e] the state of mind of the user in the input image.” Ex. 2039, ¶ 218.

Therefore, the combination of Ellenbogen and Munro does not teach “the model configured to receive an input image showing a user and predict summary statistics describing a state of mind of the user in the input image,” as required by elements 1[c], 8[c], and 15[c].

E. Challenged Dependent Claims

As discussed above, Ellenbogen and Munro do not teach several elements of claims 1, 8, and 15. Each of the remaining challenged claims in Ground 5 depend from one of those independent claims, and Ellenbogen and Munro do not render them obvious for the same reasons.

F. A POSITA Would Not Have Been Motivated to Combine Ellenbogen and Munro in the Manner Tesla Alleges

A POSITA would not have been motivated to combine Ellenbogen and Munro in the manner Tesla and Dr. Janét allege.

As discussed above, Ellenbogen and Munro relate to different fields—Ellenbogen to object recognition in images, and Munro to natural language processing of documents. *Compare, e.g.,* Ex. 1004, [0163]-[0166] *with, e.g.,* Ex. 1006, Abstract, [0003], [0004]; *see also* Ex. 2039, ¶¶ 165-167, 200-201. As Dr.

Saber explains, the differences between the fields of Ellenbogen and Munro go far beyond just the types of inputs they consider and include fundamental differences in what they aim to accomplish, the technological processes used to process those inputs, and the form of their outputs. Ex. 2039, ¶¶ 173-179. Nothing in the Petition or Dr. Janét’s declaration acknowledges much less accounts for these fundamental differences. *See* Pet. 55-76; Ex. 1003, ¶¶ 177-240. And, these differences relate directly to Tesla’s alleged combination of Ellenbogen and Munro, because Munro’s aggregation process (which is the crux of Tesla’s proposed combination) relates specifically to an issue in natural language processing, i.e., whether a “document may be reliably annotated for use in training the natural language model.” Ex. 1006, [0140]. As Dr. Saber explains, performing such an aggregation process and excluding data in an autonomous driving context would not be practical. Ex. 2039, ¶¶ 173-179. Thus, the differing practical realities of Ellenbogen and Munro’s systems and applications necessarily teaches away from combining them in the manner Tesla suggests. *Id.*, ¶¶ 173-179, 198.

Therefore, the lack of any motivation to combine Cox and Munro is an independent reason that the challenged claims are not unpatentable under Ground 5.

XI. TESLA FAILED TO SHOW THAT CERTAIN LIMITATIONS ARE PRINTED MATTER NOT ENTITLED TO PATENTABLE WEIGHT

Tesla asserts that the challenged claims include limitations that are not entitled to patentable weight under the printed matter doctrine.⁶ To prevail on this assertion, Tesla must make two showings. First, Tesla must show that the claim limitations constitute printed matter. Second, even if the claim limitations recite printed matter, Tesla must show that the limitations are *not* functionally or structurally related to the physical substrate holding the printed matter. Only if Tesla makes both showings are the claim limitations not given patentable weight in the invalidity analysis. As set forth below, Tesla has not, and cannot, meet its burden on either prong of the printed matter test. *See FanDuel, Inc. v. Interactive Games LLC*, 966 F.3d 1334, 1341 (Fed. Cir. 2020) (“In an inter partes review, the burden of persuasion is on the petitioner to prove ‘unpatentability by a preponderance of the evidence,’ 35 U.S.C. § 316(e), and that burden never shifts to the patentee.”); *Lowry*, 32 F.3d at 1584 (“Even assuming, arguendo, that data objects and data structures are analogous to printed matter . . . the burden of establishing the absence of a novel, nonobvious functional relationship rests with the PTO.”).

⁶ The printed matter doctrine is disfavored, “stand[ing] on questionable legal and logical footing.” *In re Lowry*, 32 F.3d 1579, 1583 (Fed. Cir. 1994) (citing *In re Gulack*, 703 F.2d 1381, 1385 n.8 (Fed. Cir. 1983)).

A. The Claim Limitations At Issue Are Not Printed Matter

A claim limitation is printed matter “*only* if it claims the content of information.” *In re DiStefano*, 808 F.3d 845, 848 (Fed. Cir. 2015) (emphasis added). Stated differently, “printed matter must be matter claimed *for what it communicates*.” *Id.* at 850 (emphasis added). To be considered printed matter, the claimed information must be “useful and intelligible only to the human mind.” *Lowry*, 32 F.3d at 1583 (citations omitted). “[P]rinted matter cases have no factual relevance where ‘the invention as defined by the claims *requires* that the information be processed not by the mind but by [a] computer.’” *Id.* (emphasis in original).

Cases finding printed matter illustrate these requirements. Limitations covering markings on the face of a game die were found to be printed matter. *In re Marco Guldenaar Holding B.V.*, 911 F.3d 1157, 1161 (Fed. Cir. 2018). This is because “[e]ach die’s marking or lack of marking communicates information to participants indicating whether the player has won or lost a wager” *Id.* Limitations covering markings on a vascular access port were found to be printed matter. *C R Bard Inc. v. Angiodynamics, Inc.*, 979 F.3d 1372, 1381 (Fed. Cir. 2020). The markings’ purpose was to inform a medical provider that the vascular access port was suitable for a particular type of procedures, and thus “directed to *the content of the information conveyed*.” *Id.* at 1381-82 (emphasis added). Limitations covering “identifying marks” on meat were found to be printed matter. *In re McKee*, 64 F.2d

379, 379-80 (CCPA 1933). The markings’ purpose was to inform a person of the origin of the meat or its producer. *Id.* Limitations covering a prescription label were found to be printed matter. *King Pharms., Inc. v. Eon Labs., Inc.*, 616 F.3d 1267, 1279 (Fed. Cir. 2010). The label’s instructions advised a patient to take a prescription with food. *Id.* The common thread in these cases is that the printed matter communicated specific information to a human.

Cases refusing to find a claim limitation to be printed matter are also instructive. Limitations covering a data structure including a plurality of attribute data objects stored in memory were not found to be printed matter. *Lowry*, 32 F.3d at 1583. The limitations at issue did not “claim merely the information content of a memory.” *Id.* Moreover, the Federal Circuit refused to extend the printed matter doctrine to information stored in memory, which was processed by a computer rather than conveyed to a human being. *Id.* A limitation of “selecting a first element from a database including web assets authored by third party authors and web assets provided to the user interface from outside the user interface by the user” was not found to be printed matter. *DiStefano*, 808 F.3d 848-851. The Federal Circuit found that the limitation did not claim the *content* of information. *Id.* at 851 (“Although the selected web assets can and likely do communicate some information, *the content of the information is not claimed.*”) (emphasis added).

Here, Tesla identifies nine claim limitations and alleges that they constitute printed matter. Pet. 76-77. Tesla's allegations are conclusory and lack any meaningful analysis. Tesla never explains how any of those limitations claim the content of information, as required by Federal Circuit precedent. For example, Tesla generically asserts that "each [limitation] recites a data element [] and then defines that data element based on its content." *Id.* at 77. But this allegation is plainly false as none of the limitations claim content. None of the limitations recite "matter claimed for what it communicates." *DiStefano*, 808 F.3d at 850.

For example, the first claim limitation identified by Tesla is an "image displaying one or more users." While the claimed "image" is defined as "displaying one or more users," the claim limitation does not attempt to claim specific content. The modifier "displaying one or more users" merely describes the *types* of images stored and used in the computer-implemented method. The second limitation "request to answer a question about a state of mind of the user" is similar. Nothing in this limitation defines or claims specific content of the "request to answer a question." The modifier "about a state of mind of the user" merely describes the *type* of question that will be asked.

The remaining limitations are similar in that they do not describe or claim specific content. Instead, they merely describe the types of information involved in the computer-implemented method. These limitations are very different from those

found to constitute printed matter for their communicative content, such as markings on a die that inform a game participant if they won, markings on a medical device that indicate to a doctor the device’s suitability for a particular procedure, markings on meat to identify the meat’s origin, or instructions on a prescription label to take a drug with food. *See supra*. Simply put, the limitations here bear no resemblance to those found to constitute printed matter.

Tesla also failed to address that several claim elements recite information that is not “useful and intelligible only to the human mind.” *Lowry*, 32 F.3d at 1583 (citations omitted). For example, Tesla claims that “summary statistics” are printed matter. But claim 1 is a “computer-implemented method” that involves “generating summary statistics,” “storing summary statistics . . . as part of the training data,” and training a model using the “training data.” The claimed summary statistics are never communicated to a person. *See id.* (rejecting application of printed matter doctrine to information processed by a computer). Here, the claimed information is used by a computer system to perform computational functions, not by a human reader to receive a message.

Accordingly, Tesla has not met its threshold burden of showing that the claim limitations recite printed matter.

B. The Alleged Printed Matter Is Functionally Related to the Substrate

Even if Tesla were to prevail on its claim that the limitations constitute printed matter, it cannot prevail on the second prong of the printed matter test. Printed matter is given patentable weight if the “claimed informational content has a functional or structural relationship to the substrate.” *DiStefano*, 808 F.3d at 850. A limitation deemed to be printed matter is entitled to patentable weight if it is “interrelated with the rest of the claim.” *Praxair Distrib., Inc. v. Mallinckrodt Hosp. Prods. IP Ltd.*, 890 F.3d 1024, 1032, 1035 (Fed. Cir. 2018) (“By interrelating the claimed information regarding correlations between nitric oxide, LVD, and pulmonary edema with the concrete step of discontinuing treatment of the information . . . the printed matter . . . has a functional relationship to the rest of the claim[,] giving the printed matter patentable weight.”). The Federal Circuit further “consider[s] whether the printed matter merely informs people of the claimed information, or whether it instead interacts with the other elements of the claim to create a new functionality in a claimed device or to cause a specific action in a claimed process.” *C R Bard Inc.*, 979 F.3d at 1381.

Cases finding that printed matter is entitled to patentable weight illustrate the requisite functional relationship. In *In re Miller*, the court reversed a finding of printed matter regarding “measuring indicia” and a “legend” printed on the side of a volumetric measuring device. 418 F.2d 1392, 1395-96 (CCPA 1969). These

limitations were given patentable weight because the indicia and legend enabled a user to measure specific volumes accurately and were thus integral to the functionality of the measuring device. *Id.* In *Gulack*, a claim limitation capturing a sequence of digits (“plurality of individual digits”) printed on an endless wrist band (“educational and recreational mathematical device”) constituted printed matter. 703 F.2d 1381, 1385 (Fed. Cir. 1983). Nevertheless, “the sequence deserved patentable weight because the informational content of the sequence (what numbers were represented) was functionally related to the endless-band physical structure of the substrate.” *DiStefano*, 803 F.3d at 850 (citing *Gulack*, 703 F.2d at 1385).

Here, Tesla has not, and cannot, demonstrate that the cited limitations are *not* functionally related to the substrate. Tesla asserts, in conclusory fashion, that “[t]he above printed matter limitations do not have any such functional relationship with the ‘substrate,’ which in this case is computer memory.”⁷ Pet 78. But this allegation

⁷ Tesla’s fails to explain why the alleged substrate is “computer memory.” Assuming Tesla is correct, this is reason enough to disregard Tesla’s printed matter challenge. *See Lowry*, 32 F.3d at 1583 (“[T]he Board erroneously extended a printed matter rejection under sections 102 and 103 to a new field in this case, which involves information stored in a memory. . . . The printed matter cases have no factual

is not true at all. For example, claim 1 is a “computer-implemented method” where “a plurality of images” are used to generate “training data.” The “training data” is sent to “a plurality of human observers,” whose responses are used to “generat[e] summary statistics.” The “summary statistics” are used to train an AI “model.” The “model” is then “execut[ed]” to “predict the state of mind of a user” in a new image. Each of the challenged limitations participates directly in this sequence of operations. The information is not merely stored in memory or presented to a user; it drives the generation of training data, the training of the model, and the ultimate prediction produced by the claimed method. The challenged limitations are therefore functionally intertwined with the operation of the claimed process. One would be hard pressed to find limitations more fundamentally related to the function of the computer-implemented method. Indeed, removing the challenged limitations would fundamentally alter how the claimed method operates. That is the hallmark of a functional relationship under Federal Circuit precedent.

In sum, the challenged claim limitations do not comprise printed matter. Even if they did, any alleged printed matter is still entitled to patentable weight.

relevance where ‘the invention as defined by the claims requires that the information be processed not by the mind but by a machine, the computer.’”) (citations omitted).

XII. CONCLUSION

For the foregoing reasons, Patent Owner respectfully requests that the Board find all challenged claims of the '046 patent not unpatentable.

Dated: June 15, 2026

By: / C. Austin Ginnings /
C. Austin Ginnings
Registration No. 62,706
NELSON BUMGARDNER
CONROY P.C.
745 McClintock Drive, Suite 340
Burr Ridge, Illinois 60527
Telephone: (708) 675-1583
Email: austin@nelbum.com

Attorney for Patent Owner

CERTIFICATE OF COMPLIANCE WITH WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), I certify that this Patent Owner's Response complies with the type-volume limits of 37 C.F.R. § 42.24(b)(1) because it contains 13,956 words, excluding the parts of this Patent Owner's Response that are exempted by 37 C.F.R. § 42.24(a), according to the word processing system used to prepare this Patent Owner's Response.

Dated: June 15, 2026

By: /C. Austin Ginnings/
C. Austin Ginnings
Registration No. 62,706
NELSON BUMGARDNER
CONROY P.C.
745 McClintock Drive, Suite 340
Burr Ridge, Illinois 60527
Telephone: (708) 675-1583
Email: austin@nelbum.com

Attorney for Patent Owner

CERTIFICATE OF SERVICE

The undersigned certifies that the foregoing Patent Owner's Response was served on June 15, 2026, to Lead and Back-up Counsel for Petitioner at the service address provided in Petitioner's Mandatory Notices:

Email: Roger.fulghum@bakerbotts.com
 Mark.speegle@bakerbotts.com
 Ellyar.barazesh@bakerbotts.com
 Lance.goodman@bakerbotts.com
 William.gainers@bakerbotts.com
 afawzy@tesla.com

By: / C. Austin Ginnings /
 C. Austin Ginnings
 Registration No. 62,706
 NELSON BUMGARDNER
 CONROY P.C.
 745 McClintock Drive, Suite 340
 Burr Ridge, Illinois 60527
 Telephone: (708) 675-1583
 Email: austin@nelbum.com

Attorney for Patent Owner