

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

TESLA, INC.

Petitioner,

v.

PERCEPTIVE AUTOMATA LLC

Patent Owner

IPR2025-01577

U.S. Patent No.: 11,520,346 B2

**DECLARATION OF DR. JASON JANÉT
REGARDING U.S. PATENT NO. 11,520,346**

TABLE OF CONTENTS

	Page
I. INTRODUCTION	2
A. Engagement	2
B. Background and Qualifications	2
C. Compensation and Prior Testimony	12
D. Information Considered.....	12
II. LEGAL STANDARDS FOR PATENTABILITY	13
III. SUMMARY OF THE '346 PATENT	20
A. Prosecution History	20
B. Technology Background	22
C. The Alleged Invention.....	23
IV. LEVEL OF ORDINARY SKILL IN THE ART	27
V. CLAIM CONSTRUCTION.....	27
VI. THE ASSERTED GROUNDS OF INVALIDITY	27
A. Ground 1: Claims 1-23 are unpatentable under 35 U.S.C. § 102 or 35 U.S.C. § 103 over Djuric.....	27
1. Summary of Prior Art: Djuric-EX1004	27
2. Claim 1	30
a. 1[Pre]: A method comprising:.....	30
b. 1[a]: receiving, by an autonomous vehicle, sensor data from sensors mounted on the autonomous vehicle, the sensor data comprising one or more images;.....	30
c. 1[b]: generating a point cloud representation of the surroundings of the autonomous vehicle based on the sensor data;.....	31

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

- d. 1[c]: identifying, one or more traffic entities based on the sensor data, the traffic entities representing non-stationary objects in traffic in which the autonomous vehicle is driving; 33
 - e. 1[d-i]: for each of the one or more traffic entities: determining one or more motion parameters describing movement of the traffic entity;..... 33
 - f. 1[d-ii]: [for each of the one or more traffic entities:] providing an image of the one or more images, the image showing the traffic entity as input to a machine learning model configured to receive an input image showing an input traffic entity and output summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image; 34
 - g. 1[d-iii]: [for each of the one or more traffic entities:] determining the hidden context of the traffic entity based on the output of the machine learning based model; 38
 - h. 1[d-iv]: [for each of the one or more traffic entities:] determining a region of the point cloud where the traffic entity is expected to reach within a threshold time interval; and 40
 - i. 1[d-v]: [for each of the one or more traffic entities:] modifying the region based on the hidden context of the traffic entity; and..... 43
 - j. 1[e]: navigating the autonomous vehicle so that the autonomous vehicle stays at least a threshold distance away from the modified region of each of the one or more traffic entities. 44
3. Claim 2: The method of claim 1, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in the direction having a component along the motion vector, extending the region along the direction of the motion vector.....46
4. Claim 3: The method of claim 1, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in a direction having a component opposite to direction of the motion

	vector, decreasing a size of the region along the direction of the motion vector.	48
5.	Claim 4: The method of claim 1, wherein the hidden context represents a state of mind of a user represented by the traffic entity.....	51
6.	Claim 5: The method of claim 1, wherein the hidden context represents a task that a user represented by the traffic entity is planning on accomplishing.	52
7.	Claim 6: The method of claim 1, wherein the hidden context represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity.....	53
8.	Claim 7: The method of claim 1, wherein the hidden context represents a goal of a user represented by the traffic entity, wherein the user expects to achieve the goal within a threshold time interval.	54
9.	Claim 8: The method of claim 1, wherein navigating the autonomous vehicle comprises: generating signals for controlling the autonomous vehicle based on the one or more motion parameters and the hidden context of each of the one or more traffic entities; and sending the generated signals to controls of the autonomous vehicle.	56
10.	Claim 9: The method of claim 1, wherein the sensor data represents one or more images captured by a camera mounted on the autonomous vehicle or a lidar scan captured by a lidar mounted on the autonomous vehicle.	58
11.	Claim 10.....	58
	a. 10[Pre]: A non-transitory computer readable storage medium storing instructions, that when executed by a processor, cause the processor to perform steps comprising:.....	58
	b. 10[a]: receiving, by an autonomous vehicle, sensor data from sensors mounted on the autonomous vehicle, the sensor data comprising one or more images;.....	59

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

- c. 10[b]: generating a point cloud representation of the surroundings of the autonomous vehicle based on the sensor data;..... 59
 - d. 10[c]: identifying, one or more traffic entities based on the sensor data, the traffic entities representing non-stationary objects in traffic in which the autonomous vehicle is driving; 59
 - e. 10[d-i]: for each of the one or more traffic entities determining one or more motion parameters describing movement of the traffic entity;..... 59
 - f. 10[d-ii]:[for each of the one or more traffic entities:] providing an image of the one or more images, the image showing the traffic entity as input to a machine learning model configured to receive an input image showing an input traffic entity and output summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image; 59
 - g. 10[d-iii]:[for each of the one or more traffic entities:] determining the hidden context of the traffic entity based on the output of the machine learning based model; 60
 - h. 10[d-iv]: [for each of the one or more traffic entities:] determining a region of the point cloud where the traffic entity is expected to reach within a threshold time interval; and 60
 - i. 10[d-v]: [for each of the one or more traffic entities:] modifying the region based on the hidden context of the traffic entity; and..... 60
 - j. 10[e]: navigating the autonomous vehicle so that the autonomous vehicle stays at least a threshold distance away from the modified region of each of the one or more traffic entities. 60
12. Claim 11: The non-transitory computer readable storage medium of claim 10, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in the direction having a component along the motion vector, extending the region along the direction of the motion vector. 61
13. Claim 12: The non-transitory computer readable storage medium of claim 10, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity

- represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in a direction having a component opposite to direction of the motion vector, decreasing a size of the region along the direction of the motion vector. 61
14. Claim 13: The non-transitory computer readable storage medium of claim 10, wherein the hidden context represents a state of mind of a user represented by the traffic entity. 61
15. Claim 14: The non-transitory computer readable storage medium of claim 10, wherein the hidden context represents a task that a user represented by the traffic entity is planning on accomplishing. 62
16. Claim 15: The non-transitory computer readable storage medium of claim 10, wherein the hidden context represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity. 62
17. Claim 16: The non-transitory computer readable storage medium of claim 10, wherein the hidden context represents a goal of a user represented by the traffic entity, wherein the user expects to achieve the goal within a threshold time interval. 62
18. Claim 17: The non-transitory computer readable storage medium of claim 10, wherein navigating the autonomous vehicle comprises: generating signals for controlling the autonomous vehicle based on the one or more motion parameters and the hidden context of each of the one or more traffic entities; and sending the generated signals to controls of the autonomous vehicle. 62
19. Claim 18: The non-transitory computer readable storage medium of claim 10, wherein the sensor data represents one or more images captured by a camera mounted on the autonomous vehicle or a lidar scan captured by a lidar mounted on the autonomous vehicle. 63
20. Claim 19.....63

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

- a. 19[Pre]: A computer system comprising: a processor; and a non-transitory computer readable storage medium storing instructions that when executed by the processor, cause the processor to perform steps comprising: 63
 - b. 19[a]: receiving, by an autonomous vehicle, sensor data from sensors mounted on the autonomous vehicle, the sensor data comprising one or more images; 63
 - c. 19[b]: generating a point cloud representation of the surroundings of the autonomous vehicle based on the sensor data; 63
 - d. 19[c]: identifying, one or more traffic entities based on the sensor data, the traffic entities representing non-stationary objects in traffic in which the autonomous vehicle is driving; 63
 - e. 19[d-i]: for each of the one or more traffic entities: determining one or more motion parameters describing movement of the traffic entity 64
 - f. 19 [d-ii]: [for each of the one or more traffic entities:] providing an image of the one or more images, the image showing the traffic entity as input to a machine learning model configured to receive an input image showing an input traffic entity and output summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image; 64
 - g. 19[d-iii]: [for each of the one or more traffic entities:] determining the hidden context of the traffic entity based on the output of the machine learning based model; 64
 - h. 19[d-iv]: [for each of the one or more traffic entities:] determining a region of the point cloud where the traffic entity is expected to reach within a threshold time interval; and 64
 - i. 19[d-v]: [for each of the one or more traffic entities:] modifying the region based on the hidden context of the traffic entity; and 64
 - j. 19[e]: navigating the autonomous vehicle so that the autonomous vehicle stays at least a threshold distance away from the modified region of each of the one or more traffic entities. 65
21. Claim 20: The computer system of claim 19, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in the

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

	direction having a component along the motion vector, extending the region along the direction of the motion vector.....	65
22.	Claim 21: The computer system of claim 19, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in a direction having a component opposite to direction of the motion vector, decreasing a size of the region along the direction of the motion vector.	65
23.	Claim 22: The computer system of claim 19, wherein the hidden context represents a state of mind of a user represented by the traffic entity.....	66
24.	Claim 23: The computer system of claim 19, wherein the hidden context represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity.	66
B.	Ground 2: Claims 1-23 are unpatentable under 35 U.S.C. § 103 over Djuric and Cox	66
1.	Summary of Prior Art: Cox-EX1006.....	66
2.	Motivation to Combine Djuric and Cox	68
3.	Claim 1	71
a.	1[d-ii].....	72
4.	Claims 10 and 19.....	76
5.	Claims 2-9, 11-18, and 20-23	76
C.	Grounds 3 and 4: Claims 1-23 are unpatentable under 35 U.S.C. § 103 over Djuric and Zhu or Djuric, Zhu, and Cox.....	77
1.	Summary of Prior Art: Zhu-EX1005	77
2.	Motivation to Combine Zhu with Djuric, or Djuric and Cox...	78

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

3.	Claim 1	80
a.	1[b]	81
b.	1[d-iv]	83
4.	Claims 10 and 19.....	84
5.	Claims 2-9, 11-18 and 20-23	85
D.	Grounds 5, 6, 7, and 8: Claims 1-23 are unpatentable under 35 U.S.C. § 103 over Djuric and Ross (Ground 5), Djuric, Cox, and Ross (Ground 6), Djuric, Zhu, and Ross (Ground 7), or Djuric, Cox, Zhu, and Ross (Ground 8).....	85
1.	Summary of Prior Art: Ross-EX1015.....	86
2.	Motivation to Combine Djuric, Zhu, Cox, and Ross.....	87
3.	Claims 1-3, 5, 7-12, 14, and 16-21	88
4.	Claim 4.....	89
5.	Claim 6.....	91
6.	Claim 13.....	92
7.	Claim 15.....	92
8.	Claim 22.....	92
9.	Claim 23.....	93
E.	Printed Matter.....	93
VII.	CONCLUSION	95

I. INTRODUCTION

A. Engagement

1. My name is Jason Janét, I have been retained by counsel for Tesla, Inc. (“Petitioner” or “Tesla”) as an expert witness in the above-captioned proceeding. I have been asked to provide my opinion about the state of the art of the technology described in U.S. Patent No. 11,520,346 (“the ’346 Patent”) (EX1001) and on the patentability of the claims of this Patent. The following is my written declaration on these topics.

B. Background and Qualifications

2. My credentials are set forth in full detail in my Curriculum Vitae which is attached as Exhibit A, incorporated here by reference, and summarized as follows: I am currently CTO for Applied Research Associates (ARA). I also hold the rank of Adjunct Associate Professor at Duke University and at North Carolina State University.

3. I received a Bachelor of Science in Mechanical Engineering from the University of Virginia in 1990; a Master of Integrated Manufacturing Systems from the Integrated Manufacturing Systems Engineering Institute at North Carolina State University in 1994; and a Ph.D. in Electrical and Computer Engineering from North Carolina State University in 1998.

4. Since 1991, I have been active in the robotics, sensors, artificial intelligence, pattern analysis, and automation fields. My MS thesis, entitled *Global*

Motion Planning and Self-Referencing for Autonomous Mobile Robots, leveraged graph theory and time-of-flight and/or Doppler sensors to optimize mobile robot movement through, and perception of, environments. My PhD dissertation, entitled *Pattern Analysis, Tracking and Control for Autonomous Vehicles*, leveraged multiple UGVs, time-of-flight and/or Doppler sensors, point clouds, and neural networks to enhance simultaneous localization and mapping (SLAM), as well as the sharing of knowledge between multiple different UGVs. I have authored numerous publications and have co-authored a textbook entitled *Computational Intelligence*.

5. I have designed, built, and marketed toys, robots, automated systems, measuring systems and components thereof, including ground mobile robots, unmanned aerial vehicles (UAV), unattended sensors, automated storage and retrieval systems (ASRS), submersible mobile robots, proof-of-concept extraterrestrial robots, manually-manipulated (“twiddled”) submersible foils, wall climbing mobile robots, corner-grip ballgame goals, proximity and ranging systems, point-cloud systems, pattern analysis tools, and augmented reality systems.

6. I have taught the following courses: Introduction to Robotics and Automation (Duke and NCSU); Introduction to Control Theory (Duke and NCSU); Distributed Real-Time Controls (NCSU); and myriad independent studies courses in the areas of robotics, automation, artificial intelligence, autonomy, pattern analysis, sensors, and control systems. I have also served on several MS- and PhD-

level graduate student committees, designed qualifying exam problems, served on the NCSU IMSEI Board, and participated in curriculum development for undergraduate and graduate level programs.

7. I have also started, sponsored, and advised multiple champion student teams for international robot competitions including, but not limited to, the NASA/ASCE Extra-Terrestrial Robotics Competition, the DARPA Grand Challenge, the AUVSI/ONR Autonomous Underwater Vehicle Competition, and the European CLAWAR Wall-Climbing Robot competition.

8. I have initiated multiple unmanned systems projects at both academic and industry levels. Academic unmanned aerial vehicles (UAVs) include, but are not limited to, Quadcopters with Hybrid Remote and Autonomous Control, Marsupial UAVs that Deploy and Recover Unmanned Ground Vehicles, and Wall-Climbing UAVs. Additionally, the AngelFish Cross-Domain Submersible UAV, a DARPA ASW program that I incubated and procured funding for at Teledyne, included a partnership with North Carolina State University. Unmanned underwater vehicles (UUVs) include, but are not limited to, the MicroHunter-class of very small underwater robots, the PilotFish, Gamera, Artemis and other holonomic robots, and submersible crawlers that are used in nuclear boiler water reactors, the LBC hull inspection system, and the HULLBUG hull cleaning system. Unmanned ground

vehicles (UGVs) include, but are not limited to, the miniature PERRY, wallclimbers, all-terrain and subterranean crawlers.

9. I have also supported multiple toy-focused ventures with companies including, but not limited to, Marvel, Hasbro, Twidco, Spin Master, Nottingham-Spirk, IDEO, ARA, and other firms. For Marvel, multiple Spiderman toys were conceived and developed and validated, including the Spiderman Wall-Climbing Car, the Corner-Climbing Spider, and the Inflatable Wall-Climbing Spiderman. For Hasbro, a proprietary corner-gripping technology was employed to place basketball goals in corners. For Twidco, a proprietary flapping foil technology was combined with multiple fish shapes (e.g., shark and clownfish) to enable users to create fishlike motion by “twiddling” a connecting cable from above. For Spin Master and multiple design firms, multiple wall-climbing toys were developed for small wall-climbing robot concepts including simple radio control vehicles and the “Hangtime 540” Tony Hawk aerial skateboard toy. For ARA, augmented reality systems for indoor and outdoor games.

10. In 1999, while I was employed by Nekton Research (“Nekton”), I captured and managed multiple programs sponsored by the Department of Defense (DoD) and private-sector companies that focused primarily on autonomous underwater vehicles (AUV), remotely operated vehicles (ROV), TwiddleFish toys, and indirect-fire projectiles. During my tenure, Nekton entered into a joint venture

agreement with the founders of what eventually became known as Parata Systems – a pharmacy automation solutions company, which I supported launching.

11. After leaving Nekton in March 2002, I joined a then New Jersey-based company called Avionic Instruments, Inc (“Avionic”). While at Avionic, I continued developing and marketing robots and supported engineering related to various design, manufacturing, quality and assembly issues on the core aerospace product lines. Avionic product lines include, but are not limited to, ducted fans, transformer-rectifier units (TRU), regulated TRU (RTRU), auxiliary power unit (APU) control systems, power distribution systems (PDU), frequency converters, corner clamps, and VRAM attractors/thrusters. Customers included DoD, NASA, Boeing, Sikorsky, Augusta-Westland, Dassault, and Lockheed-Martin.

12. At Avionic, I was also tasked with ruggedizing, optimizing, and commercializing two emerging proprietary technologies: the “CFG” and “VRAM”. The CFG, short for Corner Friction Grip, was a proprietary technology based on non-marring and non-adhesive materials that provided optimal grip patterns on inside and outside corners. CFG formed the basis of a company called “Shelf Works Technologies,” which was spun out of Avionic instruments in 2004, and whose products were sold in Home Depot, Bed Bath and Beyond, etc. CFG also formed the basis of multiple toys including a miniature basketball goal for Hasbro, and a corner climbing Spiderman toy for Marvel.

13. The VRAM, short for Vortex Regenerative Air/Aqueous Movement is a patented method for producing attraction forces between an object and a surface. VRAM is euphemistically referred to as a “tornado in a cup” and generally comprises both cylindrical and toroidal flow patterns that result in vacuum-like regions. The VRAM had several applications, including but not limited to: holding breaching charges and sensors against vertical or inverted surfaces; acting as an attractor for wall and ceiling climbing robots; acting as an attractor for submersible hull-crawling robots; filter-less vacuuming (conceptually similar to the Dyson vacuum); and robotic pick-and-place end-effector (conceptually similar to a gripper) to move articles from one location (e.g., a conveyor) to another (e.g., a collator). The VRAM also formed the basis for multiple toys including: a wall-climbing car for Spin Master, IDEO and Nottingham-Spirk; a Spiderman car for Marvel; an inflatable wall-climbing Spiderman for Marvel; and the Tony Hawk Hangtime 540 aerial skateboard for IDEO.

14. My work with the filter-less vacuum VRAM prompted me to explore a concept related to pharmacy automation, which ultimately led to a Ferris-wheel concept for rapidly dispensing pills. In mid-2003, my team conceived of two prescription pill counting mechanisms that actually benefitted from centrifugal forces and optimally exploited gravity and vacuum: an inner ring with vacuum-based apertures and an inner-bowl with vacuum-based apertures. The inner-bowl with

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

vacuum-based apertures was deemed most viable, was awarded two US patents, and to this day forms the basis of the RxMedic ADS™ pharmacy robot. The VRAM formed the basis of Vortex HC technologies, many of which have been licensed out to entities including, but not limited to Teledyne SeaBotix, SeaRobotics, and HDT. Both RxMedic and Vortex HC were officially spun out of Avionic in July 2004, at the date Avionic was acquired by Transdigm.

15. In late 2004, soon after Avionic was acquired, I, along with others from Avionic, secured funding to develop an alpha-level multi-dispenser pharmacy robotic system and develop the RxMedic business plan and raise multiple rounds of venture capital. In late 2006, RxMedic (called “APDS” until November 2006) was launched as a stand-alone, sole-focus venture. After the operational launch of RxMedic, I served as General Manager and eventually Chief Technical Officer. Through my roles at RxMedic, I oversaw the development of the RxMedic ADS robot, managed the intellectual property portfolio, coordinated sales and marketing, and provided strategic, fiscal, and operational leadership. In May 2010, J.M Smith Corporation acquired RxMedic, and to assist in the change of ownership, I served as a Director of RxMedic until May 2011.

16. Also in late 2004, after Avionic was acquired, I, along with others from Avionic, secured funding for Vortex HC through DoD contracts and robot sales, to continue developing the VRAM Mobile Robot Platform (VMRP – a wall-climbing

robot), the ARTEMIS AUV (a holonomic submersible robot for counter-mine and counter-obstacle operations), the submersible crawler, the nuclear-grade boiler water reactor (BWR) inspection robot, and other robot products centered around the VRAM. Some DoD programs were/are classified, for which I maintained a SECRET clearance at both the personal and facilities level, and served as the facilities security officer (FSO). In late 2006, corresponding to the full launch of RxMedic, Vortex HC technologies were largely licensed to Teledyne SeaBotix, SeaRobotics, and HDT. However, I have continued to support Vortex HC licensees and customers to-date.

17. After the sale of RxMedic, and after fulfilling my 12-month employment obligation, I joined Teledyne Technologies in Summer 2011. I served as the Senior Manager for the RTP division of Teledyne Scientific, and supported multiple DoD-sponsored robotic-focused programs. Some of these programs were/are classified, for which I maintained a personal SECRET clearance. Among these programs, were cargo unmanned ground vehicles (CUGV); squad-level autonomous unmanned ground vehicles (UGV); autonomous underwater vehicles (AUV); unmanned underwater vehicles (UUV); a cross-domain autonomous vehicle capable of transitioning between air-, surface- and underwater-domains; the EXACTO program for optically tracking and steering bullets to targets; and LIDAR systems for UGVs, as well as the mapping of urban and subterranean environments (along with sister-company Optech). Cross-domain vehicles that were evaluated,

designed and prototyped included, but were not limited to, the AngelFish (later called “EagleRay”) cross-domain submersible UAV for anti-submarine warfare, as well as a ball-shaped robot for beach-, ground- and surf-zone countermine operations. Sensor design, refinement and signal processing was a major component of each program. Sensors employed include, but are not limited to: proximity sensors; ranging sensors; electro-optical imaging; LIDAR. long-, short- and midwave infrared (IR); inertial measurement units (IMU); optical flow; and radiofrequency (RF). Additionally, control systems were designed, refined and integrated into the aforementioned systems. Most control systems were closed-loop, in that they utilized sensor-based feedback; others were open-loop, where states were estimated with little or no feedback.

18. In late 2013, Avionic, a Transdigm business unit at that point, requested that I return to turn around both a supply-chain issue, and the Sikorsky S97 Raider/JMR helicopter programs. The S97 Raider was designed to be the fastest, most maneuverable helicopter, due to its coaxial, counter-rotating variable-pitch wings, and an aft-based push-propeller. Avionic also controlled two business units named Acme Aerospace (Acme) and Aerospace Cooling Solutions (ACS). Avionic, Acme and ACS supported the S97 program, which ultimately met milestones and continues to produce multiple successful demonstrations. In 2013, Transdigm expanded my role to include directorship of the Avionic, Acme and ACS revenue

and engineering teams, and to report operational and financial status at six-week intervals. Transdigm required that I move my family to New Jersey in late 2014, which influenced my decision to resign and assume the CEO role offered to me by Delta Five Systems in Raleigh, NC.

19. As CEO of Delta Five, I led a private-equity backed, hospitality-focused venture, coordinating the company's launch and strategic path including an early-stage pivot. Initially focused on back-end robotics and automation, Delta Five shifted its focus to address the rampant bed bug problem—a top priority for hoteliers—with a novel unattended sensor and trap, called the Telemetered Pest Monitoring System (TPMS), which is based on computer vision and internet of things (IoT), and has proven capable of scaling to other pests. In addition to the TPMS, Delta Five developed a novel means to mass produce a natural, unscented aggregation pheromone that, in concert with placement and heat, lured invertebrates to the TPMS. After leading Delta Five and serving on its Board for four years, I resigned to join ARA but continue serving Delta Five as a share-holder and transition agent.

20. In 2018, I joined ARA as CTO. My primary responsibilities include leading the commercialization and technology transition efforts, which require technical, business and corporate development, as well as the procurement of capital and managing matrixed rosters.

21. I have served as an expert witness on multiple cases including cases involving IPR, ITC, contract disputes, and civil litigation in District and Federal Courts. My resume lists these cases.

C. Compensation and Prior Testimony

22. I am being compensated at a rate of \$500.00 per hour for my study and testimony in this matter. My compensation is not contingent on the outcome of this matter or the specifics of my testimony.

23. My CV, attached as Exhibit 1012, lists all other cases in which, during the previous 16 years, I have testified as an expert at trial or by deposition.

D. Information Considered

24. My opinions are based on my years of education, research and experience, as well as my investigation and study of relevant materials. In forming my opinions, I have considered the materials I identify in this report and those listed in the Exhibit List at the end of this declaration.

25. I may rely upon these materials and/or additional materials to respond to arguments raised by the Patent Owner. I may also consider additional documents and information in forming any necessary opinions — including documents that may not yet have been provided to me.

26. My analysis of the materials produced in this matter is ongoing and I will continue to review any new materials as it is provided. This declaration represents only those opinions I have formed to date.

II. LEGAL STANDARDS FOR PATENTABILITY

27. In expressing my opinions and considering the subject matter of the claims of the '346 Patent, I am relying upon certain basic legal principles that have been explained to me.

28. I understand that for an invention claimed in a patent to be found patentable under 35 U.S.C. §§ 102 and 103, it must be, among other things, new and not obvious from what was known before the invention was made.

29. I understand the information that is used to evaluate whether an invention is new and not obvious is generally referred to as “prior art” and generally includes patents and printed publications (e.g., books, journal publications, articles on websites, product manuals, etc.).

30. I understand that in this proceeding the Petitioner has the burden of proving that the claims of the '346 Patent are anticipated by or obvious from the prior art by a preponderance of the evidence. I understand that “a preponderance of the evidence” is evidence sufficient to show that a fact is more likely true than it is not.

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

31. I understand that in this proceeding, the claims should be given their ordinary and accustomed meaning as understood by one of ordinary skill in the art in view of the patent and its file history. The claims after being construed in this manner are then to be compared to the information in the prior art.

32. I understand that in this proceeding, the information that may be evaluated is limited to patents and printed publications. My analysis below compares the claims to patents and printed publications that are prior art to the claims.

33. I understand that there are two ways in which prior art may render a patent claim unpatentable. First, the prior art can be shown to “anticipate” the claim. Second, the prior art can be shown to have made the claim “obvious” to a person of ordinary skill in the art. In this declaration I focus on “obviousness.” My understanding of the legal standard for “obviousness” is set forth below.

34. I understand that a claimed invention is not patentable if it would have been obvious to a person of ordinary skill in the field of the invention at the time the invention was made.

35. I understand that the obviousness standard is defined in the patent statute (35 U.S.C. § 103(a)) as follows:

36. A patent may not be obtained, though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

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. Patentability shall not be negatived by the manner in which the invention was made.

37. I understand that the following standards govern the determination of whether a claim in a patent is obvious. I have applied these standards in my evaluation of whether the claims of the '346 Patent would have been considered obvious as of the priority date of the patent.

38. I understand that to find a claim in a patent obvious, one must make certain findings regarding the claimed invention and the prior art. Specifically, I understand that the obviousness question requires consideration of four factors (although not necessarily in the following order):

- 39. The scope and content of the prior art;
- 40. The differences between the prior art and the claims at issue;
- 41. The knowledge of a person of ordinary skill in the pertinent art; and
- 42. Whatever objective factors indicating obviousness or non-obviousness may be present in any particular case.

43. In addition, I understand that the obviousness inquiry should not be done in hindsight, but must be done using the perspective of a person of ordinary skill in the relevant art as of the effective filing date of the patent claim.

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

44. I understand the objective factors indicating obviousness or non-obviousness may include: commercial success of products covered by the patent claims; a long-felt need for the invention; failed attempts by others to make the invention; copying of the invention by others in the field; unexpected results achieved by the invention; praise of the invention by those in the field; the taking of licenses under the patent by others; expressions of surprise by experts and those skilled in the art at the making of the invention; and the patentee proceeded contrary to the accepted wisdom of the prior art. I also understand that any of this evidence must be specifically connected to the invention rather than being associated with the prior art or with marketing or other efforts to promote an invention. I am not presently aware of any evidence of “objective factors” suggesting the claimed methods are not obvious, and reserve my right to address any such evidence if it is identified in the future.

45. I understand the combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results. I also understand that an example of a solution in one field of endeavor may make that solution obvious in another related field. I also understand that market demands or design considerations may prompt variations of a prior art system or process, either in the same field or a different one, and that these variations will ordinarily be considered obvious variations of what has been described in the prior art.

46. I also understand that if a person of ordinary skill can implement a predictable variation, that variation would have been considered obvious. I understand that for similar reasons, if a technique has been used to improve one device, and a Skilled Artisan would recognize that it would improve similar devices in the same way, using that technique to improve the other device would have been obvious unless its actual application yields unexpected results or challenges in implementation.

47. I understand that the obviousness analysis need not seek out precise teachings directed to the specific subject matter of the challenged claim, but instead can take account of the “ordinary innovation” and experimentation that does no more than yield predictable results, which are inferences and creative steps that a Skilled Artisan would employ.

48. I understand that sometimes it will be necessary to look to interrelated teachings of multiple patents; the effects of demands known to the design community or present in the marketplace; and the background knowledge possessed by a person having ordinary skill in the art. I understand that all these issues may be considered to determine whether there was an apparent reason to combine the known elements in the fashion claimed by the patent at issue.

49. I understand that the obviousness analysis cannot be confined by a formalistic conception of the words “teaching, suggestion, and motivation.” I

understand that in 2007, the Supreme Court issued its decision in *KSR Int'l Co. v. Teleflex, Inc.*, 550 U.S. 398 (2007), where the Court rejected the previous requirement of a “teaching, suggestion, or motivation to combine” known elements of prior art for purposes of an obviousness analysis as a precondition for finding obviousness. It is my understanding that *KSR* confirms that any motivation that would have been known to a person of skill in the art, including common sense, or derived from the nature of the problem to be solved, is sufficient to explain why references would have been combined.

50. I understand that a person of ordinary skill attempting to solve a problem will not be led only to those elements of prior art designed to solve the same problem. I understand that under the *KSR* standard, steps suggested by common sense are important and should be considered. Common sense teaches that familiar items may have obvious uses beyond their primary purposes, and in many cases a person of ordinary skill will be able to fit the teachings of multiple patents together like pieces of a puzzle. As such, the prior art considered can be directed to any need or problem known in the field of endeavor as of the priority date of the patent and can provide a reason for combining the elements of the prior art in the manner claimed. In other words, the prior art does not need to be directed towards solving the same problem that is addressed in the patent. Further, the individual prior art references themselves need not all be directed towards solving the same problem.

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

51. I understand that obviousness does not require that the features of a secondary reference be bodily incorporated into the structure of the primary reference. Rather, the test is what the combined teachings of those references would have suggested to a Skilled Artisan. The disclosures of the prior art references need not be physically combinable. Combining the teachings of references should be the focus of the analysis.

52. I understand that an invention that might be considered an obvious variation or modification of the prior art may be considered non-obvious if one or more prior art references discourages or lead away from the line of inquiry disclosed in the reference(s). A reference does not “teach away” from an invention simply because the reference suggests that another embodiment of the invention is better or preferred. My understanding of the doctrine of teaching away requires a clear indication that the combination should not be attempted (*e.g.*, because it would not work or explicit statements saying the combination should not be made).

53. I understand that a person of ordinary skill is also a person of ordinary creativity.

54. I further understand that in many fields, it may be that there is little discussion of obvious techniques or combination, and it often may be the case that market demand, rather than scientific literature or knowledge, will drive design trends. When there is such a design need or market pressure to solve a problem and

there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within their technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense. In that instance the fact that a combination was obvious to try might show that it was obvious. The fact that a particular combination of prior art elements was “obvious to try” may indicate that the combination was obvious even if no one attempted the combination. If the combination was obvious to try (regardless of whether it was actually tried) or leads to anticipated success, then it is likely the result of ordinary skill and common sense rather than innovation.

55. I understand that in *Praxair Distribution, Inc. v. Mallinckrodt Hosp. Prods. IP Ltd.*, the Federal Circuit held that claim limitations that are directed to the content of information and that have no functional relationship are not patent eligible under 35 U.S.C. § 101. 890 F.3d 1024, 1032 (Fed. Cir. 2018) (“Claim limitations directed to the content of information and lacking a requisite functional relationship are not entitled to patentable weight because such information is not patent eligible subject matter under 35 U.S.C. § 101.”).

III. SUMMARY OF THE '346 PATENT

A. Prosecution History

56. It is my understanding that the Office issued a first non-final office action that provisionally rejected claims 1, 4-10, and 13-19 on the ground of non-

statutory double patenting over co-pending U.S. Patent Application No. 16/777,386. EX1002, 126-129. The Office also rejected independent claims 1-3, 8-12, and 17-20 under 35 U.S.C. § 103 as being unpatentable over U.S. Patent Application Publication No. 2020/0193829 (“Cheng”) in view of U.S. Patent Application Publication No. 2017/0318360 (“Tran”). *Id.*, 130-138. In response, the applicant amended claim 1 as follows, and additional independent claims 10 and 19 similarly:

“A method comprising:
receiving, by an autonomous vehicle, sensor data from sensors mounted on the autonomous vehicle, the sensor data comprising one or more images;
generating a point cloud representation of the surroundings of the autonomous vehicle based on the sensor data;
identifying, one or more traffic entities based on the sensor data, the traffic entities representing non-stationary objects in traffic in which the autonomous vehicle is driving;
for each of the one or more traffic entities:
determining one or more motion parameters describing movement of the traffic entity;
providing an image of the one or more images, the image showing the traffic entity as input to a machine learning model configured to receive an input image showing a traffic entity and output summary statistics of expected human responses describing a hidden context of the traffic entity shown in the input image;
determining ~~[[a]]~~ the hidden context of the traffic entity based on ~~[[a]]~~ the output of the machine learning based model, ~~wherein the machine learning based model is trained based on feedback received from users responsive to presenting images or videos showing traffic scenarios to the users~~;
determining a region of the point cloud where the traffic entity is expected to reach within a threshold time interval;
modifying the region based on the hidden context of the traffic entity;
navigating the autonomous vehicle so that the autonomous vehicle stays at least a threshold distance away from the modified region of each of the one or more traffic entities.

EX1002, 186-193.

57. In response to the applicant's amendments and related arguments, the Office issued a final office action that maintained the double patenting rejection, but withdrew the § 103 rejection. EX1002, 211-215. The applicant thereafter submitted a terminal disclaimer and made other minor amendments. *Id.*, 237-247, 253-254. The Office then issued a notice of allowance without commenting on the claims. *Id.*, 260-264.

58. However, as demonstrated by the grounds of this petition, the added "providing..." claim limitation used to distinguish the claims was well known, as were the Challenged Claims as a whole. These claims thus should not have been allowed, and the '346 Patent should not have issued.

B. Technology Background

59. In my opinion, the '346 Patent generally relates to machine learning models, which are computer programs that analyze input data and generate an output based on the data. Machine learning models may be trained using labeled training data that is known to represent a correct output. Human annotators (also referred to as "labelers," "reviewers," or "annotators") can be used to label such training data. After training, the model can be applied to new, non-training data to provide an output for the new data that reflects the human annotators' decision-making when labeling training data.

60. For example, it is my opinion that, to develop a vehicle detection machine learning model, human annotators may label all instances of training data images that display a vehicle. The model is then trained using the labeled training data, and, after training, the model can detect vehicles in new, non-training data from a sensor system that may produce camera image data, LIDAR point cloud data, and RADAR data. The model's detection of vehicles reflects how the human annotators would have labeled the new images because the model was trained based on the human annotators' decision-making when labeling training data. The model output can be used by an autonomous vehicle to direct vehicle navigation.

C. The Alleged Invention

61. In my opinion, the alleged invention of the '346 Patent relates to (1) gathering sensor data, including image data, with sensors of an autonomous vehicle, to generate a point cloud representation, (2) inputting the image data into a machine learning model trained to output summary statistics of expected human responses describing "hidden context," (3) determining a region of the point cloud and modifying the region based on the hidden context, and (4) navigating the autonomous vehicle to avoid the modified region. EX1001, 2:1-57, 3:19-55, 15:57-16:36, 17:46-18:26. But this alleged invention was well-known for many years in the fields of both predictive models and autonomous vehicle control, as demonstrated by the grounds herein. *See* Section VI.

62. In my opinion, the substance of the '346 Patent is nothing new. The Patent describes the well-known process of using a sensor system to “perceive” the surroundings of an autonomous vehicle by gathering sensor data, including LIDAR scans and camera images, which was a well-documented concept at least as of 2005. EX1010, pp. 3, 151, 169-174 (describing “Robot Perception” using laser range scanners and cameras to control “cars that travel by themselves”). EX1001, 1:36-38, 3:53-55, 6:66-7:6, 7:49-50, 15:57-61; EX1004, ¶¶0022, 0047, 0049-0050, 0059; EX1005, ¶¶0004, 0012, 0015, 0026, 0041. Then, a point cloud is generated based on the sensor data, which is another well-known step. EX1001, 3:19-24, 15:61-67; EX1004, ¶0022; EX1005, Abstract, ¶¶0012, 0015, 0026, 0032, 0041. The sensor data can be used to identify nearby moving objects and their associated motion parameters, yet another well-known step. EX1001, 2:44-47, 2:55-58, 7:7-15; EX1004, ¶0025, 0033, 0051.

63. The '346 Patent explains how a machine learning model uses image input to output summary statistics of expected human responses describing “hidden context” of vehicles and pedestrians, but this seemingly complex language is just jargon; “hidden context” may simply represent the behavior of other vehicles and pedestrians. EX1001, Abstract, 2:3-5. The machine learning model may be trained with training data obtained and labeled by human annotators, who provide their understanding or prediction of how traffic entities will behave, but, in my opinion,

this concept was well-known. EX1001; EX1004, ¶¶0053; Indeed, the inventor's own prior art shows that the '346 Patent should not have issued. The Cox reference that, like the '346 Patent, lists Samuel Anthony of Cambridge, Massachusetts, as an inventor (EX1006, Cover Page), discloses using human-annotated training data to ensure that machine learning model predictions are "more consistent with the decisions of the human annotators" and "better mimic human performance." EX1006, ¶¶0004, 0032. Neither the patentee nor Mr. Anthony brought the Cox reference to the patent office's attention during prosecution of the '346 Patent. *See* EX1001; EX1002.

64. After training on the labeled data, the machine learning model in the '346 Patent is used to make a prediction based on "live" input data. EX1001, 10:15-28, 13:26-48. In my opinion, this was also well-known at least as early as 2005, and the application of this step to autonomous vehicles does not save it. EX1010, pp. xvii, 3 (describing "perception and control" of "robotic systems include[ing] ... cars that travel by themselves"). EX1004, ¶¶0030-0032, 0061-0068, 0100-0102, 0111-0112; EX1005, ¶¶0073, 0079-0081 (describing inputting current images into a model that predicts future movement of objects depicted in the images).

65. The '346 Patent describes determining and modifying a region of the point cloud that a traffic entity will reach based on the hidden context. EX1001, 3:30-35, 17:46-18:12, 22:14-23:12. In my opinion, this is merely the well-known

step of predicting an object's trajectory based on its expected behavior. EX1004, ¶¶0021, 0026, 0081-0094. For example, it was well-known to generate multiple predictions of different locations an object may reach and then modify the locations by selecting one most likely location. EX1004, ¶¶0034, 0093-0094, 0112-0114.

66. The '346 Patent further describes controlling an autonomous vehicle to stay a threshold distance away from the modified region. EX1001, 3:48-52, 18:13-26, 22:41-51. However, it is my opinion that it was well known to control an autonomous vehicle to avoid the predicted future locations of objects. EX1004, ¶¶0008, 0021-0024, 0035-0036, 0050, 0078, 0095, 0103-0104, 0114, Fig. 5; EX1010, pp. xvii, 3 (explaining a goal of robotic perception and control of autonomous cars as "making car accidents a notion of the past"); EX1013, pp. 10-11, 22-23, 105-149, 356-364, 452-458 (describing configuration space (C-space) and C-space-time as approaches for planning robot motion around moving obstacles); EX1014, pp. 39-71 (describing determining a free path around obstacles in configuration space).

67. Therefore, it is my opinion that the '346 Patent describes nothing more than the application of well-known machine learning concepts to the well-known field of autonomous vehicles. The grounds of this Petition demonstrate that the Challenged Claims should not have been issued and should now be canceled. *Id.*

IV. LEVEL OF ORDINARY SKILL IN THE ART

68. It is my opinion that a person having ordinary skill in the art (“POSITA”) relevant to the ’346 Patent as of January 30, 2019 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.

V. CLAIM CONSTRUCTION

69. I understand that claim terms in an IPR are construed according to their “ordinary and customary meaning” to those of skill in the art. It is my understanding that for purposes of this proceeding, Petitioner does not present any constructions.

VI. THE ASSERTED GROUNDS OF INVALIDITY

A. Ground 1: Claims 1-23 are unpatentable under 35 U.S.C. § 102 or 35 U.S.C. § 103 over Djuric.

70. It is my opinion that Djuric discloses all the same concepts claimed in the ’346 Patent, and as such Djuric renders obvious all of the Challenged Claims.

1. Summary of Prior Art: Djuric-EX1004

71. Djuric is generally directed to “predicting the motion of objects proximate to an autonomous vehicle and controlling the autonomous vehicle.”

EX1004, ¶0002. Djuric describes an autonomous vehicle with a sensor system that includes cameras and LIDAR sensors. EX1004, ¶¶0022, 0046-0047, 0049, 0059. The sensor system captures sensor data, including camera image data and LIDAR point cloud data, associated with objects in the surroundings of the autonomous vehicle, such as other vehicles and pedestrians. *Id.* A perception system can process the sensor data to determine state data for the objects, such as “current and/or past location (also referred to as position), current and/or past speed/velocity, current and/or past acceleration, current and/or past heading, [and] current and/or past orientation.” EX1004, ¶¶0025 (parentheticals in original), 0028, 0051, 0057, 0072, 0080.

72. Djuric’s autonomous vehicle generates predictions of object motion using a predictive model, such as a “machine-learned model.” EX1004, ¶¶0027, 0047, 0053, 0061-0067. Djuric’s training methodology additionally includes “training the model [] to evaluate which trajectory of a set of trajectories is closest to a ground truth,” and therefore, in my opinion, the model produces an output closest to the human-labeled trajectories of the training data. EX1004, ¶0094.

73. Specifically, after capturing the images of objects in the autonomous vehicle’s environment, the autonomous vehicle inputs the images into the neural network model. EX1004, ¶0069. Using the images, the neural network model makes a “goal-based prediction” of future trajectories of the object, generating a plurality

of predicted trajectories for the object. EX1004, ¶¶0080-0088. For example, for an object vehicle, the predicted future trajectories may be based on driving actions such as “continu[ing] according to its current lane; merg[ing] from its current lane into an adjacent lane; [or] turn[ing] from the current lane into an adjacent road portion.” EX1004, ¶0081. Additionally, the predicted trajectories each include a plurality of way-points and associated way-point confidence levels. EX1004, ¶¶0083, 0085.

74. Using the predicted trajectories, trajectory confidence levels, waypoints, and waypoint confidence levels, the autonomous vehicle generates a motion plan and the vehicle “travel[s] in accordance with this motion plan.” EX1004, ¶0114. To execute this, a “vehicle controller can translate the motion plan into instructions for the appropriate vehicle controls (e.g., acceleration control, brake control, steering control, etc.).” EX1004, ¶¶0035, 0078.

75. It is my opinion that Djuric is analogous art to the '346 Patent, from the same field of endeavor as the '346 Patent (e.g., autonomous vehicle perception systems), and is reasonably pertinent to the particular problem that the Patent was trying to solve (e.g., how to “accurately predict motion of non-stationary traffic objects” in an autonomous vehicle’s surroundings), describing a machine learning model that is trained using human observations to predict goal-based motion of objects in the environment of an autonomous vehicle. EX1004, Abstract, EX1001, 1:34-67, 2:26-43, 3:19-35, 15:4-22.

2. Claim 1

76. It is my opinion that Djuric renders obvious claim 1.


a. 1[Pre]: A method comprising:

77. It is my opinion that Djuric discloses the preamble, to the extent it is limiting. Djuric describes systems and methods that “enable an autonomous vehicle to predict the respective motion of...objects and plan the vehicle's motion accordingly” (*[a] method*). EX1004, ¶¶0020-0024, 0035-0036, 0044-0046, 0050, 0078, 0095, 0103-0104, 0114, 0118, Fig. 9 (item 120).

b. 1[a]: receiving, by an autonomous vehicle, sensor data from sensors mounted on the autonomous vehicle, the sensor data comprising one or more images;

78. It is my opinion that Djuric teaches element 1[a] because Djuric describes that a vehicle 104 includes sensors 112, such as cameras, “and/or other types of imaging capture devices” (*sensors mounted on the autonomous vehicle*) EX1004, ¶¶0046-0047, 0059. The sensors 112 further include a “LIDAR system.” EX1004, ¶¶0022, 0047, 0049, 0059. The cameras and LIDAR system generate “sensor data 118” including “image data” from the cameras that show “objects that are within the surrounding environment of the vehicle 104,” where the objects include “vehicles, pedestrians, [and] bicycles,” as well as LIDAR data in the form of a “LIDAR point cloud.” EX1004, ¶¶0022, 0032, 0046-0049, 0050-0051, 0059, Figs. 2A-2C. A POSITA would have understood that Djuric’s image data is formed by image frames; this is because image data is formed by image frames that have

been captured by a camera. *Id.*; *see id.*, ¶¶0007, 0064, 0021, 0029-0033, 0037, 0058-0074, 0079-0080, 0099-0101. The sensor data 118 is received by autonomy computing system 114 of vehicle 104 and thus, in my opinion, teaches *receiving, by an autonomous vehicle, sensor data from sensors mounted on the autonomous vehicle, the sensor data comprising one or more images*). *Id.*, ¶0050.



EX1004, Fig. 1 (annotated)

c. 1[b]: generating a point cloud representation of the surroundings of the autonomous vehicle based on the sensor data;

79. It is my opinion that Djuric teaches element 1[b] because, as discussed for element 1[a], Djuric explains that “sensor data 118” is captured by one or more sensors and can include “image data, RADAR data, LIDAR data” including “point

cloud(s).” EX1004, ¶¶0022, 0032, 0046-0051, 0059, Figs. 2A-2C. *See* Section VI.A.2.b (element 1[a]). In my opinion, a POSITA would have understood that the LIDAR data captured by vehicle 104’s LIDAR system would have been used to generate the point clouds of the vehicle’s surroundings, thus teaching *generating a point cloud representation...based on the sensor data. Id.* This is because point clouds, which are images that display the depth/3D shape of an environment, were known as being generated using the data formed by reflected light received by a LIDAR sensor system. *Id.* EX1005, Abstract, ¶¶0004, 0015, 0026, 0043, 0048.

80. Djuric further teaches that the sensor data is “associated with the surrounding environment of the vehicle.” EX1004, ¶0022; The “LIDAR point cloud(s) and/or other data” are “associated with one or more object(s) that are proximate to the autonomous vehicle” such as “within a field of view of the sensor(s)” as well as “one or more geographic characteristics of the geographic area” such as “curbs, lane markings, [and] sidewalks.” EX1004, ¶¶0022, 0031, 0036. Djuric thus teaches that its point cloud representation generated from LIDAR data is of the surroundings of vehicle 104 (*generating a point cloud representation of the surroundings of the autonomous vehicle based on the sensor data*). *Id.*

81. Djuric’s disclosure aligns with the ’346 Patent, which discloses a vehicle “builds a point cloud representation of the surroundings of the autonomous vehicle” based on “lidar scans.” EX1001, 15:57-66.

d. 1[c]: identifying, one or more traffic entities based on the sensor data, the traffic entities representing non-stationary objects in traffic in which the autonomous vehicle is driving;

82. It is my opinion that Djuric teaches element 1[c] because Djuric describes that its autonomy computing system 114 includes a “perception system 124” and “can identify one or more objects that are proximate to the vehicle 104 based at least in part on the sensor data 118” (*identifying, one or more traffic entities based on the sensor data*), where the objects include moving “vehicles, pedestrians, [and] bicycles” that are in traffic in a roadway where the autonomous vehicle 104 is driving (*the traffic entities representing non-stationary objects in traffic in which the autonomous vehicle is driving*). EX1004, ¶¶0022, 0025, 0047, 0050-0051, 0084, 0096, 0105, Figs. 2A-2C. The objects are non-stationary objects at least because they are associated with a non-zero “current...speed/velocity” and “current...acceleration.” *Id.*, ¶¶0025, 0051, 0084, 0096, 0105.

e. 1[d-i]: for each of the one or more traffic entities: determining one or more motion parameters describing movement of the traffic entity;

83. It is my opinion that Djuric teaches element 1[d-i] because Djuric discloses that the perception system 124 processes “the sensor data 118 to detect one or more objects that are within the surrounding environment of the vehicle 104 as well as state data 130 associated therewith.” EX1004, ¶¶0021, 0051. The state data indicates “at least a current or past state of an object that is within the surrounding

environment of the vehicle 104” and “for each object” describes an estimate of the object’s current “speed/velocity” and “acceleration,” which are *motion parameters*. Therefore, in my opinion, Djuric teaches *for each of the one or more traffic entities: determining one or more motion parameters describing movement of the traffic entity*. *Id.*, ¶¶0021, 0051.

f. 1[d-ii]: [for each of the one or more traffic entities:] providing an image of the one or more images, the image showing the traffic entity as input to a machine learning model configured to receive an input image showing an input traffic entity and output summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image;

84. It is my opinion that Djuric teaches element 1[d-ii] for the following reasons.

85. First, Djuric’s vehicle 104 includes cameras that capture image data that includes “objects that are within the surrounding environment of the vehicle 104,” where the objects include moving “vehicles, pedestrians, [and] bicycles” that are in traffic in a roadway where the autonomous vehicle 104 is driving. EX1004, ¶¶0022, 0025, 0030-0032, 0046-0051, 0080-0084, 0096, 0105, Figs. 2A-2C. The image data “can include a plurality of images 306A-C (e.g., rasterized images) that include the object.” *Id.*, ¶¶0007, 0064, 0021, 0029-0033, 0037, 0058-0074, 0079-0080, 0099-0101. In my opinion, a POSITA would have understood that the image data is formed by image frames of the surrounding environment; this is because image data

is formed by image frames that have been captured by a camera. *Id.* The image data including the object is provided as an input to Djuric's model (*[for each of the one or more traffic entities:] providing an image of the one or more images, the image showing the traffic entity as input to a machine learning model*), and the model receives the image data so is thus configured to receive such data (*machine learning model configured to receive an input image showing an input traffic entity*).

86. Djuric teaches using hidden context because Djuric's model performs "a goal-based prediction." EX1004, ¶¶0080-0084. "Such an approach can allow the model 136 to consider a 'goal' that an object may aim to accomplish." *Id.*, ¶0081. Djuric's disclosure of goal-based predicted trajectories aligns with the '346 Patent's description of hidden context, which states, for example, that "[t]he hidden context may represent a goal of a user represented by the traffic entity" and "[t]he hidden context may represent a task that a user represented by the traffic entity is planning on accomplishing." EX1001, 2:11-15, 2:26-30. The model in Djuric also predicts "how an object is likely to move relative to the characteristics of the geographic area" by generating "one or more predicted future locations of the object (e.g., a predicted object trajectory), which the autonomous vehicle can use to plan and control its motion (e.g., to avoid the object)." EX1004., ¶¶0021 (parentheticals in original), 0027, 0053, 0080-0084.

87. Djuric also teaches outputting associated confidence levels for each trajectory that are each a “per-trajectory probability, quantifying how likely an object is to follow a certain trajectory.” EX1004, ¶¶0032-0034, 0080-0084, 0089-0094. Because the confidence levels are a collection of numerical measurements summarizing how likely the trajectories are, a POSITA would have recognized that these values form *summary statistics*. Djuric’s model is thus configured to output predicted trajectories and associated confidence levels that summarize the likelihoods of the trajectories (*machine learning model configured to...output summary statistics*).

88. Djuric further teaches that the output predicted trajectories and associated confidence levels (*output summary statistics*) are *of expected human responses describing a hidden context of the input traffic entity shown in the input image*. Djuric’s model is trained using “supervised training techniques” including the use of “labeled” input training data. EX1004, ¶¶0027, 0053, 0134-0136. It would have been obvious for a POSITA to use, in supervised model training, training data having object trajectories that are labeled by a human; this is because it was well-known to use human labeled training data for model training, and Djuric even suggests using “manually” labeled (i.e., human labeled) training data. *Id.*

89. Moreover, during model training, a loss function is “optimized” to “train the network to output [the] plurality of trajectories.” EX1004, ¶0092. Loss

functions were well known in the art as being used to determine the error between a model output and the labeled training data, and ultimately used to train the model to more accurately reflect labeled training data. EX1004, ¶¶0027, 0053, 0085-0094, 0134-0136. And additionally, Djuric discloses that training its model includes evaluating “which trajectory of a set of trajectories is closest to a ground truth.” EX1004, ¶0094. A POSITA would have understood that such a “ground truth,” in the context of a supervised trained model using labeled data, reflects a correct object trajectory that a human should label in an image when presented with the image. This is because ground truth is the true correct label for a particular training image.

90. It is therefore my opinion that from the aforementioned training techniques disclosed in Djuric, a POSITA would have understood that Djuric’s model is trained such that the predicted trajectories and associated confidence levels align with the expected responses from human labelers given the same input image. EX1004, ¶¶0021, 0027, 0053, 0079, 0085-0103, 0134-0136. Due to this training, the confidence levels (*output summary statistics*) for the predicted trajectories of objects in image data, as well as the predicted trajectories themselves, thus reflect expected responses from human labelers to the same input image data (*of expected human responses*).


91. And as explained above, the responses (*expected human responses*) would have been *describing a hidden context of the input traffic entity shown in the*

input image because they would reflect the potential goal of the object. EX1004, ¶¶0021, 0027, 0053, 0079, 0080-0103, 0134-0136. This is because Djuric’s model performs “a goal-based prediction” that considers the goals that an object (vehicle, bicycle, or pedestrian) in image data may “aim to accomplish.” As such, Djuric’s predicted trajectories are each representative of these goals, e.g., how an object may wish to travel in the environment around Djuric’s autonomous vehicle. A POSITA would thus have understood that the training data used to train Djuric’s model would have included labels relating to such goals.

g. 1[d-iii]: [for each of the one or more traffic entities:] determining the hidden context of the traffic entity based on the output of the machine learning based model;

92. It is my opinion that Djuric teaches element 1[d-iii]. Djuric’s model outputs a plurality of predicted trajectories for an object and associated trajectory confidence levels that correspond to *hidden context* of the object. Section VI.A.2.f (element 1[d-ii]). Because Djuric teaches using “goal-based prediction[s],” the plurality of predicted trajectories reflects one or more goals that the object may intend to pursue. *See* EX1004, ¶0081. Djuric’s system further evaluates the potential goals of a user when assigning confidence values to the predicted trajectories (*[for each of the one or more traffic entities:] determining the hidden context of the traffic entity based on the output of the machine learning based model*). EX1004, ¶¶0006, 0027, 0034, 0053, 0080-0094, 0112-0113, Figs. 8-9.

93. Djuric provides an example of this in relation to Figure 8, where a first predicted trajectory (805A) for an object vehicle predicts that the vehicle will perform a right-hand turn at an intersection and a second predicted trajectory (805B) for the object vehicle predicts that the vehicle will drive straight through the intersection. EX1004, ¶0093, Fig. 8.




Djuric, Fig. 8 (annotated)

94. When the object vehicle is traveling in the right-hand lane but is not decelerating as it approaches the intersection, Djuric's model can be trained to determine that the confidence level for trajectory 805B (traveling straight through the intersection) is higher than the level for trajectory 805A (turning right). EX1004,

¶0093, Fig. 8. This example demonstrates that Djuric's system determines the likely goal(s) (*hidden context*) when assigning confidence levels to a predicted trajectory *[for each of the one or more traffic entities:] determining the hidden context of the traffic entity based on the output of the machine learning based model.*

h. 1[d-iv]: [for each of the one or more traffic entities:] determining a region of the point cloud where the traffic entity is expected to reach within a threshold time interval; and

95. It is my opinion that Djuric teaches element 1[d-iv] because Djuric's predicted trajectories and associated confidence values identify *a region of the point cloud where the traffic entity is expected to reach within a threshold time interval* as claimed. Section VI.A.2.f (element 1[d-ii]). The predicted trajectories are formed of waypoints that are each "indicative of a predicted future location of the object." EX1004, ¶¶0007, 0082-0094. Figure 8 shows example future waypoint locations (annotated in red) where an object is expected to be based on the predicted trajectories (annotated in green). *Id.*, ¶¶0082-0094, Fig. 8. The image shown by Figure 8 is an "output 800 of the model 136." EX1004, ¶0083.



Djuric, Fig. 8 (annotated)

96. It is my opinion that Djuric’s predicted trajectory and waypoints teach *determining a region [...] where the traffic entity is expected to reach within a threshold time interval*. EX1004, ¶¶0007, 0022, 0026, 0071, 0082-0094, Fig. 8. For example, Djuric explains that each “predicted trajectory can include a plurality of predicted future locations of an object at one or a plurality of future times.” EX1004, ¶0071; *see also id.*, ¶¶0026, 0052, 0072-0073, 0102.

97. These “predicted future locations” of each “predicted trajectory” all define a determined *region [...] where the traffic entity is expected to reach* because

they are the areas in the image of the environment that Djuric's model predicts the object will potentially travel to. EX1004, ¶¶0026, 0052, 0071-0073, 0102. Moreover, the "plurality of future times" teaches the *threshold time interval* because they form an interval of time within which the object is predicted to reach the last waypoint of the predicted trajectory. *Id.* Additionally, Djuric's "vehicle computing system has limited computational resources onboard the vehicle and is running in real-time, as the vehicle is travelling." EX1004, ¶0037. It is my opinion that it would have been obvious to a POSITA that the limited computational resources limit the time horizon of Djuric's predictions. Any prediction Djuric makes of a region an object may reach occurs within the computational window or time horizon (*threshold time*) defined by Djuric's limited computational resources. Thus, the trajectories and waypoints further teach *a region [...] where the traffic entity is expected to reach within a threshold time interval* because they define the areas in the environment that Djuric's model predicts the object will potentially travel to within the computational window or time horizon defined by Djuric's limited computational resources.

98. In my opinion, a POSITA would further have found it obvious that the trajectories and waypoints of objects in the environment image around Djuric's autonomous vehicle teaches "*a region of the point cloud.*" The sensor data used in Djuric includes "LIDAR point cloud(s) [...] associated with one or more object(s)

[...] [and] one or more geographic characteristics of the geographic area.” EX1004, ¶0022. Thus, a POSITA would have understood that the surroundings mapped in Djuric (e.g., as shown in Figure 8) and displayed in the image of the environment is LIDAR point cloud data captured by sensors. *Id.* Thus, the specific trajectories and waypoints output from the model reflect future locations of the object in the LIDAR point cloud image (*[for each of the one or more traffic entities:] determining a region of the point cloud where the traffic entity is expected to reach within a threshold time interval*). EX1004, ¶¶0007, 0022, 0026, 0052, 0059, 0071, 0082-0094, Fig. 8; See Section VI.A.2.c (element 1[b]).

i. 1[d-v]: [for each of the one or more traffic entities:] modifying the region based on the hidden context of the traffic entity; and

99. It is my opinion that Djuric teaches element 1[d-v]. As explained for element 1[d-iv], the collection of predicted trajectories in Djuric defines a *region of the point cloud where the traffic entity is expected to reach within a threshold time interval*. See Section VI.A.2.h. Djuric teaches that this *region* is modified because Djuric’s system “determine[s] a selected predicted trajectory” from “among the plurality of predicted trajectories of the object” by selecting the determined trajectory that has the highest confidence level. EX1004, ¶¶0006, 0034, 0080-0094, 0113, Fig. 9. Selecting a one “winner” trajectory out of the initial plurality involves eliminating some locations that the object was predicted to reach on other

trajectories, thereby *modifying the region* that the object is predicted to reach. *Id.* Moreover, the plurality of predicted trajectories for the object and associated trajectory confidence levels correspond to the *hidden context* of the object. Sections VI.A.2.f (element 1[d-ii]), VI.A.2.g (element 1[d-iii]). Accordingly, by selecting the determined predicted trajectory for each object (and by extension, discarding other trajectories for the object) based on the output confidence levels, Djuric teaches *[for each of the one or more traffic entities:] modifying the region based on the hidden context of the traffic entity*. *Id.*, ¶¶0085-0094, 0113, Fig. 8.

j. 1[e]: navigating the autonomous vehicle so that the autonomous vehicle stays at least a threshold distance away from the modified region of each of the one or more traffic entities.

100. Djuric teaches element 1[e] because Djuric describes that after its model “output[s] one or more predicted future locations of the object (e.g., a predicted object trajectory),” the autonomous vehicle uses the trajectory to “plan and control its motion (e.g., to avoid the object).” EX1004, ¶¶0021 (parentheticals in original), 0034-0035, 0052, 0071, 0077-0078, 0103, 0112-0114; Sections VI.A.2.h, VI.A.2.i (element 1[d-iv], element 1[d-v]). Specifically, “the vehicle computing system 102 can control a motion of the vehicle 104 based at least in part on the selected predicted trajectory.” EX1004, ¶¶0113-0114. Djuric’s system “best navigates the vehicle relative to the objects at their future locations.” *Id.*, ¶¶0035, 0077. Djuric also lists “enhancing passenger/vehicle/object safety” and “improving

vehicle efficiency by reducing collisions” as benefits of its system. *Id.*, ¶¶0021, 0036. In my opinion, a POSITA would therefore have understood that the vehicle is controlled such that it stays a distance greater than zero away from surrounding traffic (*navigating the autonomous vehicle so that the autonomous vehicle stays at least a threshold distance away from the modified region of each of the one or more traffic entities*). EX1004, ¶¶0021, 0077, 0103, 0112-0114, Figs. 6, 9.

101. Additionally, it is my opinion that a POSITA would have found it obvious to control a vehicle such that it stays a distance greater than zero away from surrounding traffic and the selected predicted trajectory for an object (*navigating the autonomous vehicle so that the autonomous vehicle stays at least a threshold distance away from the modified region of each of the one or more traffic entities*) based on Djuric’s goals of “enhancing passenger/vehicle safety” and “reducing collisions.” EX1004, ¶¶0021, 0036, 0077, 0103, 0112-0114, Figs. 6, 9; EX1013, pp. 10-11, 22-23, 105-149, 356-364, 452-458; EX1014, pp. 39-71. This would have been the predictable modification of well-known prior art elements (e.g., Djuric’s autonomous vehicle control based on object trajectory prediction; the well-known method of not causing collisions by staying away from surrounding traffic) according to known methods to yield predictable results (Djuric’s model avoiding predicted trajectories to reduce collisions). In my opinion, a POSITA would have also had a reasonable expectation of success because Djuric already explains that its


system is designed to ensure that its autonomous vehicle does not collide with objects in its environment.

3. Claim 2: The method of claim 1, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in the direction having a component along the motion vector, extending the region along the direction of the motion vector.

102. Djuric renders obvious claim 2. Djuric teaches that the “objects” being analyzed “can include, for example, other vehicles, pedestrians, bicycles, etc.” (*the traffic entity represents a user*). EX1004, ¶0022.

103. Djuric’s model outputs predicted trajectories and waypoints (*the region*) based on an input “combined data set” that includes “state data” of the object in image data (*the region is in a direction determined based on a motion vector of the traffic*). EX1004, ¶¶0025, 0079-0080, 0084, 0107-0109; *see* Section VI.A.2.h (element 1[d-iv]). The state data for an object may include “an estimate of the object's current and/or past location (also referred to as position), current and/or past speed/velocity, current and/or past acceleration, current and/or past heading, current and/or past orientation.” EX1004, ¶¶0025, 0051. At least the “speed/velocity” and “heading” combine to form a vector, and thus it is my opinion that the output waypoints are *based on a motion vector of the traffic entity*.

104. As discussed above, Djuric's model selects a predicted trajectory (which includes waypoints) based on a determined hidden context of the object (*modifying the region based on the hidden context comprises: responsive to determining based on the hidden context...*). EX1004, ¶¶0081, 0085-0088, 0094, 0113; *see* Section VI.A.2.i (element 1[d-v]). Djuric describes an example scenario where the model determines that a vehicle, a *traffic entity represented by a user*, traveling in a right-hand travel lane toward an intersection “may be likely to take a right-hand turn or travel straight through the intersection.” EX1004, ¶0084.



Djuric Fig. 8 (In part, annotated)


105. As Figure 8 shows, if the model selects the predicted trajectory corresponding to “travel[ing] straight through the intersection” (left trajectory, above), then the predicted trajectory and waypoints extend straight. EX1004, ¶0084, Fig. 8. In my opinion, Djuric therefore teaches *responsive to determining[...]*that the user represented by the traffic entity is likely to move “straight through the intersection” (in the direction having a component along the motion vector), extending the predicted trajectories and waypoints (the region) in that same direction (along the direction of the motion vector). EX1004, Fig. 8.

4. Claim 3: The method of claim 1, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in a direction having a component opposite to direction of the motion vector, decreasing a size of the region along the direction of the motion vector.

106. Djuric renders obvious Claim 3. Claim 3 is similar to Claim 2 except that, instead of the user being likely to move in the same direction as its current motion vector, in Claim 3 the user is likely to move in an opposite direction. *Id*; see *supra* Section VI.A.3. So, for the same reasons as in Claim 2, it is my opinion that Djuric discloses “wherein the region is in a direction determined based on a motion vector of the traffic entity,” “wherein the traffic entity represents a user,” and

“modifying the region based on the hidden context.” Supra Section VI.A.3 (Claim 2); EX1004, ¶¶0025, 0079-0080, 0084, 0107-0109; see also Section VI.A.2.h (element 1[d-iv]), VI.A.2.i (element 1[d-v]).

107. Similar to Claim 2, Djuric also discloses or renders obvious *decreasing a size of the region along the direction of the motion vector* in response to *determining based on the hidden context that the user represented by the traffic entity is likely to move in a direction having a component opposite to direction of the motion vector*. Djuric’s example scenario where the model determines that a vehicle “may be likely to take a right-hand turn or travel straight through the intersection” is again informative. EX1004, ¶0084.



Djuric Fig. 8 (In part, annotated)

108. As shown in the above the right-side trajectory, when the model determines that the vehicle is likely to turn right (instead of going straight) the predicted trajectories and waypoints do not extend as far in the direction that the vehicle is currently traveling (*decreasing a size of the region along the direction of the motion vector*). EX1004, ¶0084, Fig. 8. It is my opinion that a POSITA would have understood that the right turn trajectory involves acceleration opposite in a direction relative to how the vehicle was traveling (e.g., braking), which is *a component [of the predicted motion] opposite to direction of the motion vector*.

109. Additionally, it would have been obvious that a right-hand turn could be on a road that forms an acute angle with the current path, such that the motion vector after the turn has a component opposite to the original path of the car. For example, it would have been obvious to a POSITA that the vehicle could make a U-turn or a second right hand turn to fully travel in the opposite direction. Consistent with Figure 8, it would have been obvious to a POSITA that such a turn would further reduce the extent to which the predicted trajectories and waypoints (the *region*) extend in the direction of the vehicle's current movement (*decreasing a size of the region along the direction of the motion vector*).

110. It is also my opinion that a POSITA would have found it obvious that by selecting a right-hand turn trajectory with a motion vector that has a component opposite to the original path of the car, and by discarding the left-side trajectory of FIG. 8 that goes "straight through the intersection," *a size of the region along the direction of the motion vector is decreas[ed]*. EX1004, ¶¶0085-0094, Fig. 8.

5. Claim 4: The method of claim 1, wherein the hidden context represents a state of mind of a user represented by the traffic entity.

111. Djuric renders obvious claim 4. As explained for claim 1, Djuric teaches that the predicted trajectories for an object—for example, a vehicle—represent a *hidden context* for the vehicle because they represent a "goal" the vehicle "may aim to accomplish" such as turning, merging, or continuing straight. EX1004,


¶¶0032-0043, 0081-0085, 0107-0113; Section VI.A.2.f (element 1[d-ii]). Because this *hidden context* represents the goals/intentions of a vehicle (and in turn, the vehicle's driver), the *hidden context* represents the state of mind as to where the vehicle's driver wants to go (*represents a state of mind of a user represented by the traffic entity*). EX1004, ¶¶0032-0043, 0081-0085, 0107-0113. This aligns with the '346 Patent, which explains that a state of mind may include "that a pedestrian is likely to cross the street or that an oncoming vehicle is unlikely to be willing to yield." EX1001, 12:7-12.

6. Claim 5: The method of claim 1, wherein the hidden context represents a task that a user represented by the traffic entity is planning on accomplishing.

112. Djuric renders obvious claim 5. As explained for claim 1, Djuric teaches that the predicted trajectories for an object—for example, a vehicle—represent a *hidden context* for the vehicle because they represent goals the vehicle "may aim to accomplish" including tasks like turning, merging, or continuing straight. EX1004, ¶¶0032-0043, 0081-0085, 0107-0113; Section VI.A.2.f (element 1[d-ii]). Because this *hidden context* represents the goals/intentions of a vehicle (and in turn, the vehicle's driver), the *hidden context* represents a task that the driver is planning on accomplishing (*represents a task that a user represented by the traffic entity is planning on accomplishing*). EX1004, ¶¶0032-0043, 0081-0085, 0107-0113.

7. Claim 6: The method of claim 1, wherein the hidden context represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity.

113. Djuric renders obvious claim 6. As explained for claim 1, Djuric teaches that the predicted trajectories for an object—for example, a vehicle—represent a *hidden context* for the vehicle because they represent goals the vehicle “may aim to accomplish” such as turning, merging, or continuing straight. EX1004, ¶¶0032-0043, 0081-0085, 0107-0113; Section VI.A.2.f (element 1[d-ii]). Because this *hidden context* represents the goals/intentions of a vehicle (and in turn, the vehicle’s driver), it is my opinion that it would have been obvious to a POSITA that the *hidden context* can represent that driver’s awareness of the autonomous vehicle (*represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity*). EX1004, ¶¶0032-0043, 0081-0085, 0107-0113. For example, a POSITA would have understood, in my opinion, that another vehicle’s (e.g., object 202 in Figure 5) awareness of the autonomous vehicle 104 would impact “the potential risk to vehicle 104” that is relevant to vehicle 104 passing object 202. *See* EX1004, ¶0076. This also aligns with the disclosure in Djuric that the model considers the awareness of other surroundings, like a pedestrian’s awareness of an obstruction in a sidewalk. EX1004, ¶¶0074-0077.




EX1004, Fig. 5

8. Claim 7: The method of claim 1, wherein the hidden context represents a goal of a user represented by the traffic entity, wherein the user expects to achieve the goal within a threshold time interval.

114. Djuric renders obvious claim 7. As explained for claim 1, Djuric teaches that the predicted trajectories for an object—for example, a pedestrian—represent a *hidden context* for the vehicle because they represent goals the pedestrian “may aim to accomplish” or may expect to achieve. EX1004, ¶¶0032-0043, 0074-0077, 0081-0085, 0107-0113; Section VI.A.2.f (element 1[d-ii]). Because this *hidden context* represents the goals/intentions of a pedestrian, the *hidden context* represents a goal of a user represented by the traffic entity, wherein the user expects

to achieve the goal within a threshold time interval. EX1004, ¶¶0032-0043, 0081-0085, 0107-0113.

115. For example, Djuric discloses predicting a trajectory 504 based on a pedestrian's goal to "bypass an obstruction 312 in a sidewalk." EX1004, ¶¶0074-0077, Fig. 5.



EX1004, Fig. 5

116. Djuric's pedestrian (user) *expects to achieve the goal within a threshold time interval* because, as discussed for element 1[d-iv], the trajectory is predicted to occur within a threshold time interval. EX1004, ¶¶0007, 0022, 0026, 0071, 0082-0094, Fig. 8. For example, Djuric explains that each "predicted trajectory can include

a plurality of predicted future locations of an object at one or a plurality of future times.” EX1004, ¶0071. These “predicted future locations” of each “predicted trajectory” define a region that the object is expected to reach, and the “plurality of future times” define the corresponding *threshold time interval*. *Id.* Djuric also teaches *a threshold time interval* in an additional way. Djuric discloses that “the vehicle computing system has limited computational resources onboard the vehicle and is running in real-time, as the vehicle is travelling.” EX1004, ¶0037. A POSITA would have found it obvious that the limited computational resources thus limit the time horizon of Djuric’s predictions. Therefore, Djuric’s prediction that the pedestrian plans to bypass an obstruction on the sidewalk occurs within the computational window or time horizon defined by Djuric’s limited computational resources.

9. Claim 8: The method of claim 1, wherein navigating the autonomous vehicle comprises: generating signals for controlling the autonomous vehicle based on the one or more motion parameters and the hidden context of each of the one or more traffic entities; and sending the generated signals to controls of the autonomous vehicle.

117. Djuric renders obvious claim 8. As described for claim 1, Djuric describes a method “of predicting object motion and controlling autonomous vehicles.” EX1004, ¶¶0095, 0104; *See* Section VI.A.2.e (element 1[d-i]). Djuric further discloses that navigating the autonomous vehicle includes “send[ing] one or more **control signals** to the responsible vehicle control component (e. g., braking

control system, steering control system, acceleration control system) to execute the instructions and implement the motion plan” (*generating signals for controlling the autonomous vehicle[...] and sending the generated signals to controls of the autonomous vehicle*). EX1004, ¶0078 (parenthetical in original, emphasis added); *see also id.*, ¶¶0035, 0079, 0105, 0107.

118. Djuric further teaches that the autonomous vehicle is controlled *based on one or more motion parameters and the hidden context of each of the one or more traffic entities*. As discussed above, Djuric discloses *motion parameters* used in the model: “state data” for an object (vehicle, pedestrian, bicyclist) that “describe[s] an estimate of the object's current and/or past location (also referred to as position), current and/or past speed/velocity, current and/or past acceleration, current and/or past heading, current and/or past orientation” which are all motion (e.g., velocity, acceleration) parameters of the object (*motion parameters...of the traffic entity*). EX1004 ¶¶0025 (parenthetical in original), 0051, 0071, 0102, 0127; Section VI.A.2.e (element 1[d-i]). The state data is used as part of a “combined data set” that is input to the model to generate the predicted trajectories for the objects. EX1004, ¶¶0104-0114. The act of *controlling the autonomous vehicle* in Djuric is thus based on the predicted trajectory (*hidden context*), and also based on the state data (*motion parameters*) of the detected objects (*each of the one or more traffic entities*).

EX1004, ¶¶0021-0025, 0034-0036, 0051, 0071, 0077-0078, 0083-0094, 0102-0114, 0127, Fig. 9; VI.A.2.j (element 1[e]).

10. Claim 9: The method of claim 1, wherein the sensor data represents one or more images captured by a camera mounted on the autonomous vehicle or a lidar scan captured by a lidar mounted on the autonomous vehicle.

119. Djuric renders obvious claim 9. Djuric discloses that the autonomous vehicle includes one or more sensors, such as a LIDAR system and one or more cameras, which acquire sensor data associated with objects proximate to the autonomous vehicle. EX1004, ¶¶0022, 0047, 0051, 0064. The sensor data can include “image data” and “LIDAR data,” such as point cloud data (*the sensor data represents one or more images captured by a camera mounted on the autonomous vehicle or a lidar scan captured by a lidar mounted on the autonomous vehicle*). *Id.*

11. Claim 10

120. Djuric renders obvious claim 10 for the same reasons discussed with respect to claim 1, as indicated below. *See* Section VI.A.2

a. 10[Pre]: A non-transitory computer readable storage medium storing instructions, that when executed by a processor, cause the processor to perform steps comprising:

121. Djuric teaches element 10[Pre]. For example, Djuric discloses that “vehicle computing system 102 can include one or more computing device(s) 701. The computing device(s) 701 of the vehicle computing system 102 can include processor(s) 702 and a memory 704 (e.g., onboard the vehicle 104).” EX1004, ¶0116

(parenthetical in original). Additionally, “memory 704 can include one or more non-transitory computer-readable storage media,” (*[a] non-transitory computer readable storage medium*) and “the memory 704 can store instructions 706 that when executed by the one or more processors 702 cause the one or more processors 702 (the computing system 102) to perform operations.” (*storing instructions, that when executed by a processor, cause the processor to perform*). EX1004, ¶¶0117-0118; *See also* Section VI.A.2.a (element 1[Pre]).

b. 10[a]: receiving, by an autonomous vehicle, sensor data from sensors mounted on the autonomous vehicle, the sensor data comprising one or more images;

122. *See* Section VI.A.2.b (element 1[a]).

c. 10[b]: generating a point cloud representation of the surroundings of the autonomous vehicle based on the sensor data;

123. *See* Section VI.A.2.c (element 1[b]).

d. 10[c]: identifying, one or more traffic entities based on the sensor data, the traffic entities representing non-stationary objects in traffic in which the autonomous vehicle is driving;

124. *See* Section VI.A.2.d (element 1[c]).

e. 10[d-i]: for each of the one or more traffic entities determining one or more motion parameters describing movement of the traffic entity;

125. *See* Section VI.A.2.e (element 1[d-i]).

f. 10[d-ii]:[for each of the one or more traffic entities:] providing an image of the one or more images,

the image showing the traffic entity as input to a machine learning model configured to receive an input image showing an input traffic entity and output summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image;

126. *See* Section VI.A.2.f (element 1[d-ii]).

g. 10[d-iii]: [for each of the one or more traffic entities:] determining the hidden context of the traffic entity based on the output of the machine learning based model;

127. *See* Section VI.A.2.g (element 1[d-iii]).

h. 10[d-iv]: [for each of the one or more traffic entities:] determining a region of the point cloud where the traffic entity is expected to reach within a threshold time interval; and

128. *See* Section VI.A.2.h (element 1[d-iv]).

i. 10[d-v]: [for each of the one or more traffic entities:] modifying the region based on the hidden context of the traffic entity; and

129. *See* Section VI.A.2.i (element 1[d-v]).

j. 10[e]: navigating the autonomous vehicle so that the autonomous vehicle stays at least a threshold distance away from the modified region of each of the one or more traffic entities.

130. *See* Section VI.A.2.j (element 1[e]).

12. Claim 11: The non-transitory computer readable storage medium of claim 10, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in the direction having a component along the motion vector, extending the region along the direction of the motion vector.

131. In my opinion, Djuric renders obvious claim 11 for the same reasons discussed above in Section VI.A.3 (claim 2).

13. Claim 12: The non-transitory computer readable storage medium of claim 10, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in a direction having a component opposite to direction of the motion vector, decreasing a size of the region along the direction of the motion vector.

132. In my opinion, Djuric renders obvious claim 12 for the same reasons discussed above in Section VI.A.4 (Claim 3).

14. Claim 13: The non-transitory computer readable storage medium of claim 10, wherein the hidden context represents a state of mind of a user represented by the traffic entity.

133. In my opinion, Djuric renders obvious claim 13 for the same reasons discussed above in Section VI.A.5 (Claim 4).

15. Claim 14: The non-transitory computer readable storage medium of claim 10, wherein the hidden context represents a task that a user represented by the traffic entity is planning on accomplishing.

134. In my opinion, Djuric renders obvious claim 14 for the same reasons discussed above in Section VI.A.6 (Claim 5).

16. Claim 15: The non-transitory computer readable storage medium of claim 10, wherein the hidden context represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity.

135. In my opinion, Djuric renders obvious claim 15 for the same reasons discussed above in Section VI.A.7 (Claim 6).

17. Claim 16: The non-transitory computer readable storage medium of claim 10, wherein the hidden context represents a goal of a user represented by the traffic entity, wherein the user expects to achieve the goal within a threshold time interval.

136. In my opinion, Djuric renders obvious claim 16 for the same reasons discussed above in Section VI.A.8 (Claim 7).

18. Claim 17: The non-transitory computer readable storage medium of claim 10, wherein navigating the autonomous vehicle comprises: generating signals for controlling the autonomous vehicle based on the one or more motion parameters and the hidden context of each of the one or more traffic entities; and sending the generated signals to controls of the autonomous vehicle.

137. In my opinion, Djuric renders obvious claim 17 for the same reasons discussed above in Section VI.A.9 (Claim 8).

19. Claim 18: The non-transitory computer readable storage medium of claim 10, wherein the sensor data represents one or more images captured by a camera mounted on the autonomous vehicle or a lidar scan captured by a lidar mounted on the autonomous vehicle.

138. In my opinion, Djuric renders obvious claim 18 for the same reasons discussed above in Section VI.A.10 (Claim 9).

20. Claim 19

139. Djuric renders obvious claim 19 for the same reasons discussed with respect to claims 1 and 10, as indicated below. *See* Sections VI.A.2, VI.A.11 (Claims 1, 10).

a. 19[Pre]: A computer system comprising: a processor; and a non-transitory computer readable storage medium storing instructions that when executed by the processor, cause the processor to perform steps comprising:

140. *See* Sections VI.A.11.a (element 10[Pre]), VI.A.2.a (element 1[Pre]).

b. 19[a]: receiving, by an autonomous vehicle, sensor data from sensors mounted on the autonomous vehicle, the sensor data comprising one or more images;

141. *See* Section VI.A.2.b (element 1[a]).

c. 19[b]: generating a point cloud representation of the surroundings of the autonomous vehicle based on the sensor data;

142. *See* Section VI.A.2.c (element 1[b]).

d. 19[c]: identifying, one or more traffic entities based on the sensor data, the traffic entities

representing non-stationary objects in traffic in which the autonomous vehicle is driving;

143. *See* Section VI.A.2.d (element 1[c]).

e. 19[d-i]: for each of the one or more traffic entities: determining one or more motion parameters describing movement of the traffic entity

144. *See* Section VI.A.2.e (element 1[d-i]).

f. 19 [d-ii]: [for each of the one or more traffic entities:] providing an image of the one or more images, the image showing the traffic entity as input to a machine learning model configured to receive an input image showing an input traffic entity and output summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image;

145. *See* Section VI.A.2.f (element 1[d-ii]).

g. 19[d-iii]: [for each of the one or more traffic entities:] determining the hidden context of the traffic entity based on the output of the machine learning based model;

146. *See* Section VI.A.2.g (element 1[d-iii]).

h. 19[d-iv]: [for each of the one or more traffic entities:] determining a region of the point cloud where the traffic entity is expected to reach within a threshold time interval; and

147. *See* Section VI.A.2.h (element 1[d-iv]).

i. 19[d-v]: [for each of the one or more traffic entities:] modifying the region based on the hidden context of the traffic entity; and

148. *See* Section VI.A.2.i (element 1[d-v]).

j. 19[e]: navigating the autonomous vehicle so that the autonomous vehicle stays at least a threshold distance away from the modified region of each of the one or more traffic entities.

149. *See* Section VI.A.2.j (element 1[e]).

21. Claim 20: The computer system of claim 19, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in the direction having a component along the motion vector, extending the region along the direction of the motion vector.

150. Djuric renders obvious claim 20 for the same reasons discussed above in Section VI.A.3 (Claim 2).

22. Claim 21: The computer system of claim 19, wherein the region is in a direction determined based on a motion vector of the traffic entity, wherein the traffic entity represents a user, wherein modifying the region based on the hidden context comprises: responsive to determining based on the hidden context that the user represented by the traffic entity is likely to move in a direction having a component opposite to direction of the motion vector, decreasing a size of the region along the direction of the motion vector.

151. Djuric renders obvious claim 21 for the same reasons discussed above in Section VI.A.4 (Claim 3).

23. Claim 22: The computer system of claim 19, wherein the hidden context represents a state of mind of a user represented by the traffic entity.

152. Djuric renders obvious claim 22 for the same reasons discussed above in Section VI.A.5 (Claim 4).

24. Claim 23: The computer system of claim 19, wherein the hidden context represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity.

153. Djuric renders obvious claim 23 for the same reasons discussed above in Section VI.A.7 (Claim 6).

B. Ground 2: Claims 1-23 are unpatentable under 35 U.S.C. § 103 over Djuric and Cox

154. To the extent Djuric alone does not teach *output[ing] summary statistics of expected human responses describing a hidden context* as recited in element 1[d-ii], the claimed operations are also disclosed by Cox. It is my opinion that, for the reasons described below, it would have been obvious to a POSITA to combine these features from Cox with Djuric.

1. Summary of Prior Art: Cox-EX1006

155. Cox is a PCT patent application by Harvard College, which lists Samuel Anthony as an inventor—the same Samuel Anthony listed as an inventor of the '346 Patent. EX1001, Cover; EX1006, Cover. Cox is a material reference to the '346 Patent that teaches several concepts relevant to the Challenged Claims but was never

brought to the attention of the Patent Office during prosecution or otherwise considered. EX1001, Cover; *see generally* EX1002.

156. Cox is directed toward “machine learning augmented by measurements acquired from human populations.” EX1006, Abstract, ¶¶0001-0004. Cox teaches crowd-sourcing training data used to train machine learning algorithms. EX1006, Abstract, ¶¶0001-0004, 0015-0016, 0025-0034, 0044, 0060. Just like the ’346 Patent, Cox discloses compiling an n-amount of images (training objects), transmitting the images to a plurality of human annotators (human labelers), compiling statistical information regarding the distribution of human responses (including the actual responses and psychometric data associated with the responses), and training a supervised learning model using gathered information “consistent with the decisions of the human annotators 160.” *Cf* EX1006, ¶¶0029-0032, Figs. 1A-B, 3 *with* EX1001, 2:31-43, 8:8-10:4; Fig. 4. The goal of Cox’s machine learning techniques is to “better mimic human performance.” EX1006, ¶0004. Moreover, Cox touts that its “supervised learning” techniques may be “directly applied to several important domains where machine learning is found” and expressly suggests applying the invention to “driverless / semiautonomous automobile[]” applications. EX1006, ¶¶0056, 0060.

157. Cox is analogous art to the ’346 Patent, from the same field of endeavor as the Patent (e.g., autonomous vehicle perception systems), and is reasonably

pertinent to the particular problem that the Patent was trying to solve (e.g., how to “accurately predict motion of non-stationary traffic objects” in an autonomous vehicle’s surroundings), describing the gathering training data from a population of human labelers, training a model using the training data, and executing it to output predictions. EX1006, ¶¶0029-0032, Figs. 1A-B, 3; EX1001, 1:34-67, 2:26-43, 3:19-35, 15:4-22; Fig. 4.

2. Motivation to Combine Djuric and Cox

158. It is my opinion that a POSITA would have been motivated to combine Djuric such that Djuric’s “goal-based prediction” (*hidden context*; EX1004, ¶¶0080-0084) and trajectory confidence levels (*summary statistics*) are generated in accordance with human-annotated/ manually-labeled training data, as disclosed by Cox. EX1004, ¶0053; EX1006, ¶¶0004, 0027-0032, 0051. A POSITA would have been motivated to combine Djuric and Cox in this way for several reasons. Djuric describes that the model can be trained with “supervised training techniques” and “using labeled state data and geographic data with known future location(s) and/or trajectories of objects (e.g., as captured in a data log, manually).” EX1004, ¶¶0027, 0053, 0134. As taught by Cox, using human-annotated image data to train Djuric’s model would enable the model to provide predictions that are “more consistent with the decisions of the human annotators,” and that “better mimic human performance.” EX1006, ¶¶0004, 0027-0032, 0051.

159. For example, Cox simply expands on Djuric's disclosure of a model for an autonomous vehicle that is trained using supervised training and manual labeling. EX1004, ¶¶0027, 0053, 0134, 0137; Section VI.A.1. Djuric discloses generating training data, using the training data to train a model, and executing the model to predict trajectories for objects around the autonomous vehicle. EX1004, ¶¶0053, 0084-0086, 0090-0094, 0113, 0134-0139. Cox expands on this by identifying such a model as a supervised learning model, and teaching that the model is trained and optimized using human observations, so that model predictions better reflect human decision making. EX1006, ¶¶0007-0010, 0015-0019, 0025-0038, 0051, 0060, Fig. 4B. The goal of Cox's machine learning techniques is to "better mimic human performance," and Cox explains that its machine learning "may be applied to any form of supervised learning, including neural networks." EX1006, ¶¶0004, 0060; see also *id.*, ¶¶0008-0010, 0015-0019, 0025-0035, 0051. Cox thus provides a natural expansion and improvement on Djuric's base disclosure.

160. In my opinion, a POSITA seeking to implement Djuric's model would also have looked to Cox to provide more instructions for training a machine learning model to predict human behavior. Both are directed to machine learning technologies and involve (1) gathering training data, (2) training a model with the training data, and (3) executing the model, and as noted above, both are explicitly applicable to autonomous vehicle applications. Cox even acknowledges as much,

explaining that its disclosure is relevant to many “visual recognition problems” including those related to detecting “objects” such as a “car,” as well as “driverless” vehicle applications. EX1006, ¶¶0016, 0056. Cox therefore provides an explicit teaching, suggestion, or motivation to make a combination with a system for controlling driverless (or autonomous) vehicles, just like Djuric.

161. It is my opinion that a POSITA would have also been motivated to combine the teachings of Djuric and Cox because the combination would have benefited Djuric’s model. Enhancing Djuric’s model with Cox’s human-annotated training data that considers the state of mind of training subjects, as well as Cox’s associated model training and optimization processes, would have further refined Djuric’s model so that it better considers the state of mind of other road users and what they might be thinking when present in the autonomous vehicle’s environment. EX1006, ¶¶0007-0010, 0015-0019, 0025-0035, 0051, 0060. This would have improved the Djuric model’s goal-based predictions and thus helped reduce collisions between the autonomous vehicle and human drivers or pedestrians.

162. In my opinion, all the reasons explained above establish that a POSITA would have been motivated to combine Djuric and Cox as described. This combination would have been the predictable combination of well-known prior art elements (e.g., the Djuric model being a supervised learning model; Cox’s model training and optimization techniques) according to known methods to yield

predictable results (the Djuric-Cox model generating predictions for objects shown in image data that are “more consistent with the decisions of the human annotators” who annotated the training data). EX1006, ¶¶0032, 0037. The combination would have further used known techniques (known machine learning model training and optimization techniques as taught by Cox) to improve similar devices (the Djuric model for goal-based predictions) in the same way (using Cox’s aforementioned techniques). Using human-annotated image data in Djuric (as taught by Cox) would have been well within the skill of a POSITA, and a POSITA would have had an expectation of success in the modification because of the similarity to the training approach disclosed in Djuric.

3. Claim 1

163. Cox’s teachings only impact how a model generates *output summary statistics of expected human responses describing a hidden context* in element 1[d-ii]. The combination of Djuric and Cox teaches element 1[d-ii] for the reasons explained below, as well as the reasons addressed in Ground 1. Sections VI.A.2, VI.A.2.f (element 1[d-ii]). Moreover, the combination of Djuric and Cox teaches the remaining elements of claim 1 for the same reasons as addressed in Ground 1. *Id.* I note that where later elements (e.g., 1[d-iii]) refer back to the analysis in 1[d-ii], the additional analysis described in Ground 2 for 1[d-ii] is equally applicable to those later elements as well.

a. 1[d-ii]

164. The combination of Djuric and Cox teaches element 1[d-ii] for the same reasons addressed in Ground 1 (Section VI.A.2.f), and also in an additional way. Specifically, and as described further below, the Djuric-Cox model generates predicted trajectories (*hidden context*) and trajectory confidence levels (*summary statistics*) reflective of human-labeled image data that is used to train the model (*expected human responses*). Section VI.A.2.f (element 1[d-ii]); EX1006, ¶¶0019, 0029-0038, 0060; Fig. 4B.

165. Cox teaches that human-labeled image data is used to train a prediction model, so that the model's output includes *summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image*. EX1006, ¶¶0019, 0029-0038, 0060; Fig. 4B. For example, Cox discloses that training data is generated by a “crowd-sourc[ed]” and “large population” of human annotators that annotate a “training object” such as an image, to generate “statistical information” based on all the responses received from the multiple different human annotators. EX1006, ¶¶0007-0010, 0015-0019, 0025-0035, 0051. Cox then describes training a supervised machine learning model using the statistical information. EX1006, ¶¶0007-0010, 0015-0019, 0025-0035, 0051. “[T]he subject matter of each training object has certain features that may be classified by at least some of the human annotators 160.” *Id.*, ¶0018, Figs. 1A, 1B. For example, the

human annotator responses describe a state of mind of a subject person in the image, e.g., “an emotion that corresponds to the [sample training object] image, e.g., jealous, panicked, arrogant, or hateful.” *Id.*, ¶0016 (bracketing added); *id.*, ¶¶0007-0010, 0015-0019, 0025-0035, 0051. Cox discloses that the statistical information is “collected and applied to the training of computer-vision systems.” EX1006, ¶0027. Cox’s model can be utilized after training “to make predictions (based on various query objects) **that are more consistent with the decisions of the human annotators 160.**” *Id.*, ¶¶0027-0032 (parenthetical in original, emphasis added), 0051. The goal of Cox’s machine learning techniques is to “better mimic human performance.” *Id.*, ¶0004.

166. Cox additionally teaches that the model’s output represents *summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image* because Cox’s training data in the form of annotation responses and psychometric data from the human annotators 160 is utilized for training the model, “resulting in a human-weighted loss function 425 computationally derived by the penalization module 180.” EX1006, ¶¶0031-0032, Fig. 4B. The loss function “includes penalties for misclassification of later presented query objects, as graphically illustrated in graph 430, and the magnitude of the penalties increases with increasing deviation from the classification data received from the human annotators 160.” EX1006, ¶¶0032, 0034. That is, “[o]nce the human

responses have been analyzed to quantify dominant patterns of error, the models may be translated into human-weighted loss functions for maximum-margin kernel machines and penalties for margins not consistent with human data.” EX1006, ¶¶0032, 0034.

167. Cox also describes that its model is trained to minimize the difference between the model’s predicted output and the human-labeled training data. EX1006, ¶¶0019, 0029-0038, 0060; Fig. 4B. Cox also discloses that its model training includes “classifying, by the computer processor, at least some of the training objects based at least in part on an initial loss function,” which refers to executing the model on training data images. EX1006, ¶¶0008, 0019, 0029-0037; Fig. 4B (430). Cox then teaches “comparing” the training data to the classification data output by the model to “identify, within the training data, features misclassified in comparison to the classification data,” and “incorporating [] penalties for misclassification within the initial loss function to generate the human-weighted loss function.” *Id.* In particular, Cox discloses that the “penalties for misclassification may be assigned based at least in part on the psychometric data” and “the magnitude of the penalties increases with increasing deviation from the classification data received from the human annotators 160,” which ensures that the trained model is “more consistent with responses from the human annotators.” EX1006, ¶¶0008, 0032. Accordingly, Cox describes that its model is trained by predicting classifications of training data and optimizing the

model such that the difference between the predicted classifications output by the model and the training data is minimized. EX1006, ¶¶0019, 0029-0038, 0060; Fig. 4B.

168. Thus, as applied to Djuric, Cox teaches that the predicted trajectories and trajectory confidence levels are generated in accordance with the human-annotated/ manually-labeled training data (*summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image*). EX1004, ¶¶0080-0084; EX1006, ¶¶0004, 0027-0032, 0051; Section VI.B.2. The human-annotated/labeled training data includes annotated/labeled image data and is used to train the Djuric model to provide predictions that are “more consistent with the decisions of the human annotators,” and that “better mimic human performance.” EX1006, ¶¶0004, 0032. This means the model’s outputs are representative of how humans would have annotated or labeled images input to the model if the images were part of the training data. EX1006, ¶¶0004, 0019, 0029-0038, 0051, 0060; Fig. 4B. As such, the Djuric-Cox combination teaches that the model receives an image input and outputs predicted trajectories and trajectory confidence levels (*[for each of the one or more traffic entities:] providing an image of the one or more images, the image showing the traffic entity as input to a machine learning model configured to receive an input image showing an input traffic entity and output summary statistics*) (see Section VI.A.2.f, element 1[d-ii]), and Cox, as

applied to Djuric, teaches the predicted confidence levels (*summary statistics*) represent a likelihood of how humans would have annotated/labeled the input image data, thus teaching *summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image*. See Section VI.A.2.f (element 1[d-ii]).

4. Claims 10 and 19

169. Djuric and Cox render obvious claims 10 and 19 for the same reasons discussed above with respect to claim 1. The combination with Cox only impacts the analysis for elements 10[d-ii] and 19[d-ii]. Elements 10[d-ii] and 19[d-ii] are taught by the Djuric-Cox combination for the reasons stated above for 1[d-ii]. Section VI.B.3.a (Ground 2, element 1[d-ii]). Djuric and Cox further teach elements 10[d-ii] and 19[d-ii] as well as the remaining elements claims 10 and 19 for the same reasons addressed for Djuric alone in Ground 1. Sections VI.A.11 (claim 10), VI.A.20 (claim 19). Like for claim 1, where later elements (e.g., 10[d-iii] and 19[d-iii]) refer back to the analysis in 10[d-ii] or 19[d-ii], the additional analysis described in Ground 2 for 10[d-ii] and 19[d-ii] is equally applicable to those later elements as well. Sections VI.A.11 (claim 10), VI.A.20 (claim 19).

5. Claims 2-9, 11-18, and 20-23

170. Djuric and Cox render obvious claims 2-9, 11-18, and 20-23 for the same reasons discussed with respect to Djuric in Section VI.A above.

C. Grounds 3 and 4: Claims 1-23 are unpatentable under 35 U.S.C. § 103 over Djuric and Zhu or Djuric, Zhu, and Cox

171. To the extent Djuric alone (or the combination of Djuric and Cox) does not disclose or render obvious (1) using sensor data to generate a point cloud representation, and/or (2) determining and modifying a region of the point cloud, such operations are also disclosed by Zhu. Claims 1-23 would thus be obvious over the further combination of Zhu, with either Djuric (Ground 3) or with Djuric and Cox (Ground 4).

1. Summary of Prior Art: Zhu-EX1005

172. Zhu is generally directed toward generating a drivable road surface representation “based on data from cameras or other image generation devices and data from laser light detection and ranging (LIDAR) devices.” EX1005, ¶0004. The cameras and LIDAR devices are used by an autonomous vehicle “to accurately detect and map the drivable road surfaces proximate to or ahead of the location of the autonomous vehicle.” *Id.*, ¶0015. The autonomous vehicle can “produce accurate drivable road surface representations in a drivable road surface map,” and “use the drivable road surface map to safely and efficiently control the vehicle... while avoiding obstacles and safely controlling the vehicle.” EX1005, ¶¶0013, 0015.

173. Zhu describes using “cameras or image capturing devices [that] may include one or more devices configured to capture a plurality of images of the

environment of the vehicle.” EX1005, ¶0026. Zhu explains that LIDAR sensors generate “point cloud data” that represents “measured distances at various points in three dimensional (3D) space in front of or adjacent to a vehicle on which the...LIDAR unit is mounted.” *Id.*, ¶¶0004, 0012, 0015, 0026. Zhu teaches to “combine” lidar point cloud data and image data by “project[ing]” the “3D point cloud data onto the 2D image data to produce mapped image/point cloud data” which thereby forms a “drivable road surface map that can be used to distinguish drivable areas of a roadway from non-drivable areas.” *Id.*, ¶¶0040, 0043-0049.

174. Therefore, in my opinion, Zhu is analogous art to the '346 Patent, from the same field of endeavor as the Patent (e.g., autonomous vehicle perception systems), and is reasonably pertinent to the particular problem that the Patent was trying to solve (e.g., how to “accurately predict motion of non-stationary traffic objects” in an autonomous vehicle’s surroundings), describing a sensor system that combines lidar point cloud data and image data to generate a vehicle environment representation. EX1005, ¶¶0004, 0012, 0013, 0037; EX1001, 1:34-67, 2:26-43, 3:19-35, 15:4-22.

2. Motivation to Combine Zhu with Djuric, or Djuric and Cox

175. In my opinion, a POSITA would have been motivated to combine Djuric and Zhu (and alternatively along with Cox) such that Djuric’s camera image data and LIDAR point cloud data are combined to generate a drivable road surface

map, where the point cloud frames captured by LIDAR devices are projected onto the image data captured by cameras to generate the drivable road surface map as taught by Zhu. In my opinion, the further inclusion of Cox does not impact the motivation to combine Djuric with Zhu, or a POSITA's expectation of success in making the combination.

176. Both Djuric and Zhu are directed to autonomous vehicle vision systems and the control of such vehicles based on captured sensor data. Djuric already describes using LIDAR sensors and cameras to capture images of an autonomous vehicle's environment and detect objects within the environment. EX1004, ¶¶0022, 0032, 0046-0049, 0050-0051, 0059, Figs. 2A-2C. Zhu explains that combining the point cloud data output by LIDAR sensors and image data output by cameras allows for "accurate drivable road surface representations in a drivable road surface map." EX1005, ¶0013. This allows for an autonomous vehicle to "accurately detect and map the drivable road surfaces proximate to or ahead of the location of the autonomous vehicle." *Id.*, ¶0015.

177. It is my opinion that a POSITA would therefore have been motivated to use Zhu's teachings to further improve Djuric's object detection using LIDAR sensors and cameras. As taught by Zhu, combining the data allows objects to be detected and for drivable road surfaces (i.e., where objects are not located in a road)

to be detected. This improves the navigation of Djuric's autonomous vehicle in relation to predicted trajectories.

178. For the foregoing reasons, it is my opinion that a POSITA would have been motivated to combine Djuric and Zhu as described. This combination would have been the predictable combination of well-known prior art elements (e.g., Djuric's autonomous vehicle with sensors that capture image data and LIDAR point cloud data; Zhu's techniques for LIDAR and image data combination) according to known methods to yield predictable results (generating a drivable road surface map by combining Djuric's camera image data and LIDAR point cloud data). The combination would have further used known techniques (the combination of LIDAR point cloud data and image data) to improve similar devices (Djuric's autonomous vehicle with sensors that capture image data and LIDAR point cloud data) in the same way (using Zhu's aforementioned techniques). Lastly, because Djuric and Zhu disclose similar approaches to collecting sensor data (including LIDAR point cloud data) and making predictions, using Zhu's technique for combining LIDAR and image data in Djuric would have been well within the skill of a POSITA, and a POSITA would have had an expectation of success in the modification.

3. Claim 1

179. Zhu's teachings only impact the analysis relating to point clouds in elements 1[b] and 1[d-iv]. The combination of Djuric and Zhu (for Ground 3) and


the combination of Djuric, Cox, and Zhu (for Ground 4) teaches elements 1[b] and 1[d-iv] for the reasons explained below, as well as the reasons addressed in Ground 1. Sections VI.A.2, VI.A.2.c (element 1[b]), VI.A.2.h (element 1[d-iv]). Moreover, the combination of Djuric and Zhu (Ground 3) and the combination of Djuric, Cox, and Zhu (Ground 4) teach the remaining elements of claim 1 for the same reasons as addressed in Ground 1. *Id.* Additionally, for Ground 4, the combination of Djuric, Cox, and Zhu teaches element 1[d-ii] for the same reasons discussed in Ground 2. Section VI.B.3.a. I note that where later elements refer back to the analysis in 1[b] or 1[d-iv], the additional analysis described in Grounds 3/4 for 1[b] or 1[d-iv] is equally applicable to those later elements as well.

a. 1[b]

180. Zhu describes an autonomous vehicle having a vision system for the vehicle's surroundings that uses cameras and LIDAR. EX1005, Abstract, ¶¶0004-0015, 0021-0027, 0037-0047, Figs. 1-3. Zhu explains that LIDAR sensors generate "point cloud data" that represents "measured distances at various points in three dimensional (3D) space in front of or adjacent to a vehicle on which the...LIDAR unit is mounted." *Id.*, ¶¶0004, 0012, 0015, 0026. Zhu teaches a process to "combine" the lidar point cloud data and the image data by "project[ing]" the "3D point cloud data onto the 2D image data to produce mapped image/point cloud data,"

which thereby forms a “drivable road surface map that can be used to distinguish drivable areas of a roadway from non-drivable areas.” *Id.*, ¶¶0040, 0043-0049.

181. Thus, applying Zhu’s teachings to Djuric, the Djuric-Zhu combination would have specified that that Djuric’s camera image data and LIDAR point cloud data would have been combined to generate a drivable road surface map, where the point cloud frames captured by lidar devices are projected onto the image data captured by cameras to generate the drivable road surface map (*generating a point cloud representation of the surroundings of the autonomous vehicle based on the sensor data*). EX1005, Abstract, ¶¶0004-0015, 0021-0027, 0037-0047, Figs. 1-3; *See* Section VI.A.2.b (element 1[a]). Figure 7 below shows an example drivable road surface map. EX1005, ¶0046, Fig. 7.



EX1005, Fig. 7

b. 1[d-iv]

182. As explained in Ground 1, Djuric teaches *determining a region [...]* where the traffic entity is expected to reach within a threshold time interval by generating predicted trajectories and waypoints where a traffic entity is expected to reach in a threshold time interval. See Section VI.A.2.h. As explained below, in combination with Zhu it would have been obvious to a POSITA that this *region* would be on the drivable surface generated using LIDAR point cloud data and camera image data (*a region of the point cloud*). EX1004, ¶¶0007, 0022, 0026, 0052, 0059, 0071, 0082-0094, Fig. 8; See Section VI.C.3.a (element 1[b]).

183. Djuric explains that the image shown by Figure 8 is an “output 800 of the model 136.” EX1004, ¶0083. And Djuric explains that such an output may be formed using “image data from a front-facing camera, other cameras, RADAR data, LIDAR data, and/or data from other sensors.” *Id.*, ¶¶0022, 0059. As taught by Zhu, such LIDAR point cloud data and camera image data can be combined to form a drivable road surface map (the point cloud representation). See Section VI.C.3.a (element 1[b]). Thus, it is my opinion that, by using Djuric’s model 136 to analyze Zhu’s drivable surface map (which is formed by LIDAR point cloud data and camera image data, as contemplated by Djuric), a POSITA would have understood that Djuric’s vehicle computing system would have determined the locations of the predicted trajectories and waypoints on the drivable surface map. EX1004,

¶¶0007, 0022, 0026, 0052, 0059, 0071, 0082-0094, Fig. 8; *See* Section VI.C.3.a (element 1[b]).

184. Djuric’s predicted trajectories and waypoints would define *where the traffic entity is expected to reach within a threshold time interval* for the same reasons addressed in Ground 1. *See* Section VI.A.2.h (element 1[d-iv]). So, when plotted onto the drivable surface map as taught by Zhu, the predicted trajectories and waypoints are “*a region of the point cloud where the traffic entity is expected to reach within a threshold time interval.*”

4. Claims 10 and 19

185. Djuric and Zhu (Ground 3) and Djuric, Cox, and Zhu (Ground 4) render obvious claims 10 and 19 for the same reasons discussed above with respect to claim 1. As explained above, the combination with Zhu only impacts the analysis for elements 10[b], [d-iv], 19[b], and 19[d-iv]. Elements 10[b], [d-iv], 19[b], and 19[d-iv] are taught by the Djuric-Zhu combination and the Djuric-Cox-Zhu combination for the reasons stated above for elements 1[b] and 1[d-iv]. *Supra*, Sections VI.C.3.a and VI.C.3.b. Djuric and Zhu (for Ground 3) and Djuric, Cox, and Zhu (for Ground 4) further teach elements 10[b], 10[d-iv], 19[b], and 19[d-iv], as well as the remaining elements of claims 10 and 19, for the same reasons addressed

for Djuric alone in Ground 1.⁸ *Supra*, Sections VI.A.11 (claim 10), VI.A.20 (claim 19). Additionally, for Ground 4, the combination of Djuric, Cox, and Zhu teach elements 10[d-ii] and 19[d-ii] for the same reasons discussed in Ground 2. *Supra*, Sections VI.B.3, VI.B.4.

5. Claims 2-9, 11-18 and 20-23

186. Djuric and Zhu (Ground 3) and Djuric, Cox and Zhu (Ground 4) render obvious claims 2-9, 11-18, and 20-23 for the same reasons discussed with respect to Djuric in Section VI.A above. *See* Section VI.B.5.

D. Grounds 5, 6, 7, and 8: Claims 1-23 are unpatentable under 35 U.S.C. § 103 over Djuric and Ross (Ground 5), Djuric, Cox, and Ross (Ground 6), Djuric, Zhu, and Ross (Ground 7), or Djuric, Cox, Zhu, and Ross (Ground 8)

187. To the extent Djuric alone (or any of the Djuric-Cox, Djuric-Zhu, or Djuric-Cox-Zhu combinations) does not render obvious dependent claims 4, 6, 13, 15, 22, and 23, further combination with Ross renders these dependent claims obvious. The analysis for the remainder of the claims is unchanged from Grounds 1 through 4. Claims 1-23 would thus be obvious over the further combination of Ross,

⁸ Like for claim 1, where later elements refer back to the analysis in 10[b], 10[d-iv], 19[b] or 19[d-iv], the additional analysis described in Grounds 3/4 for 10[b], 10[d-iv], 19[b] or 19[d-iv] is equally applicable to those later elements as well.

with either Djuric (Ground 5), Djuric and Cox (Ground 6), Djuric and Cox (Ground 7), or Djuric, Cox, and Zhu (Ground 8).

1. Summary of Prior Art: Ross-EX1015

188. Ross is generally directed to techniques for “maneuvering” an autonomous vehicle. EX1015, 1:30-31, 3:39-42, 3:64-5:5. The goal of Ross is to provide an autonomous vehicle control system that provides “[r]easoning about the intents of objects, such as other road users” (e.g., other vehicles, pedestrians, bicyclists) and allow the autonomous vehicle to “function in a more ‘human-like’ or ‘polite’ way.” *Id.*, 5:6-11, 15:1-16:52. For example, Ross explains that an action of a “pedestrian waiting to cross the road” can correspond to an intent of the pedestrian “to cross the road when clear,” i.e., when there is no traffic, and it is safe for the pedestrian to cross.

189. Ross is analogous art to the ’346 Patent, from the same field of endeavor as the Patent (e.g., autonomous vehicle perception systems), and is reasonably pertinent to the particular problem that the Patent was trying to solve (e.g., how to “accurately predict motion of non-stationary traffic objects” in an autonomous vehicle’s surroundings); Ross describes a model that provides “[r]easoning about the intents of objects, such as other road users.” EX1015, 5:6-11, 15:1-16:52.

2. Motivation to Combine Djuric, Zhu, Cox, and Ross

190. In my opinion, POSITA would have been motivated to combine Djuric and Zhu and/or Cox as discussed in Sections VI.B.2 and VI.C.2. A POSITA would have been further motivated to combine Ross with Djuric (and optionally Zhu and/or Cox) such that the Djuric-Ross model generates “goal-based prediction[s]” using the Djuric’s goals and Ross’s intents or actions.

191. As discussed regarding elements 1[d-ii]-1[d-iii], Djuric describes generating “goal-based prediction[s]” of trajectories of objects in a vehicle’s environment. EX1004, ¶¶0081-0085; *See* Sections VI.A.2.f, VI.A.2.g (elements 1[d-ii], 1[d-iii]). It is my opinion that a POSITA would have understood that Djuric’s model is not limited to predicting trajectories based solely on the example goals disclosed by Djuric. Instead, a POSITA would have found it obvious to supplement Djuric’s use of goals with the prediction of intent and/or actions disclosed by Ross, which would expand the scenarios Djuric’s model is capable of predicting. Ross describes an intent model that is configured to predict a “set of possible intents or hypotheses identifying possible next actions (and predicted trajectories), a given point in time when the action is likely to occur, and associated likelihood values.” EX1015, 1:55-2:13, 4:33-46, 9:13-40, 15:1-21. Ross discloses that “each possible intent of the set includes a corresponding next action to be initiated by the road user.” *Id.*, 2:43-44. A POSITA would have found it obvious to

supplement Djuric's list of goals with the intents and/or actions disclosed by Ross and thereby expand the predictive capabilities of Djuric's model, leading to more accurate predictions.

192. For the foregoing reasons, a POSITA would have been motivated to combine Djuric and Ross (and optionally Zhu and/or Cox) as described. This combination would have been the predictable combination of well-known prior art elements (e.g., Djuric's model for goal-based predictions of object movement; Ross's techniques for intent-based predictions of road users) according to known methods to yield predictable results (using Djuric's model to predict object motion based on their intents as determined based on Ross's teachings). The combination would have further used known techniques (Ross's intent determination for predictions) to improve similar devices (Djuric's goal-based prediction model) in the same way (using Ross's aforementioned techniques). Lastly, because Djuric and Ross disclose similar approaches to generating predictions for autonomous vehicles, using Ross's technique for intent determination in Djuric would have been well within the skill of a POSITA, and a POSITA would have had an expectation of success in the modification.

3. Claims 1-3, 5, 7-12, 14, and 16-21


193. Djuric (Ground 1), Djuric and Cox (Ground 2), Djuric and Zhu (Ground 3), and Djuric, Cox, and Zhu render obvious claims 1-3, 5, 7-12, 14, and 16-21 for

the same reasons discussed above in Grounds 1 through 4. *See* Sections VI.A (Ground 1), VI.B (Ground 2), and VI.C (Grounds 3 and 4).

4. Claim 4

194. Djuric and Ross render obvious claim 4 for the same reasons discussed in Section VI.A.5. To the extent Djuric does not explicitly disclose the hidden context represents a state of mind, the combination of Djuric and Ross renders this obvious.

195. Ross describes a model that is configured to predict a “set of possible intents or hypotheses identifying possible next actions (and predicted trajectories), a given point in time when the action is likely to occur, and associated likelihood values.” EX1015, 1:55-2:13, 4:33-46, 9:13-10:10, 15:1-16:52. Such actions reflect the associated “possible intent[s]” (*hidden context*) associated with the object. For example, as shown in annotated Figure 9 below, an intent may be that the driver of vehicle 920 intends that vehicle 920 will “exit from parking lot 924 and cross edge 922,” and thereby move (shown in red) into the path of autonomous vehicle 100 (shown in green) when it is “clear.” (*the hidden context represents a state of mind of a user represented by the traffic entity*). EX1015, 16:5-41, Fig. 9.




EX1015, Fig. 9 (in part, annotated)

196. Ross's teachings align with the '346 patent, which explains that a state of mind may include "that a pedestrian is likely to cross the street or that an oncoming vehicle is unlikely to be willing to yield." EX1001, 12:7-12. Like the '346 patent, Ross describes predicting an intention (exiting a parking lot). EX1015, 16:5-41, Fig. 9.

5. Claim 6

197. Djuric and Ross render obvious claim 6 for the same reasons discussed in Section VI.A.7. To the extent Djuric does not explicitly disclose the hidden context represents a degree of awareness of the autonomous vehicle by a user, the combination of Djuric and Ross renders this obvious.

198. Another predicted intent described by Ross is that a pedestrian “would like to cross the roadway” that a vehicle 100 is traveling on. EX1015, 16:23-41, Fig. 6.



EX1015, Fig. 6 (in part, annotated)

199. As described by Ross, a pedestrian may “cross[] a roadway” on which vehicle 100 is traveling or “cross[] a roadway in a crosswalk” on the road vehicle

100 is traveling, and when the crossing traverses the path of the vehicle (shown in green above), the vehicle must stop for the pedestrian to cross. *Id.*, 4:47-65, 9:41-10:10, 16:23-41, Fig. 6. Moreover, Ross explains that an action of a “pedestrian waiting to cross the road” can correspond to an intent of the pedestrian “to cross the road when clear,” i.e., when there is no traffic, and it is safe for the pedestrian to cross. *Id.*, 15:50-16:22. Ross further explains a pedestrian “may want to cross but... won't step out in front of a vehicle unless that vehicle slows down and allows for them to cross.” *Id.*, 3:53-56. In such a case, the predicted intent (*hidden context*) thus demonstrates an awareness of the vehicle 100 driving in the roadway toward the pedestrian (*represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity*).

6. Claim 13

200. Djuric and Ross render obvious claim 13 for the same reasons discussed above in Section VI.D.4 (Claim 4).

7. Claim 15

201. Djuric and Ross render obvious claim 15 for the same reasons discussed above in Section VI.D.5 (Claim 6).

8. Claim 22

202. Djuric and Ross render obvious claim 22 for the same reasons discussed above in Section VI.D.4 (Claim 4).

9. Claim 23

203. Djuric and Ross render obvious claim 23 for the same reasons discussed above in Section VI.D.5 (Claim 6).

E. Printed Matter

204. I understand that Claim limitations directed to the content of information and lacking a requisite functional relationship are not entitled to patentable weight because such information is not patent eligible subject matter.

205. Because they recite the content of information without a functional relationship, the following limitations are printed matter:

- 1[a], 10[a], 19[c] – “the sensor data comprising one or more images”
- 1[b], 10[b], 19[d] – “a point cloud representation of the surroundings of the autonomous vehicle”
- 1[d-i], 10[d-i], 19[d-i] – “motion parameters describing movement of the traffic entity”
- 1[d-ii], 10[d-ii], 19[d-ii] – “[the]/[an input] image showing [the]/[a] traffic entity,” “summary statistics of expected human responses describing a hidden context of the input traffic entity shown in the input image”
- 1[d-iv], 10[d-iv], 19[d-iv] – “a region of the point cloud where the traffic entity is expected to reach within a threshold time interval”
- Claims 4, 13 – “hidden context represents a state of mind of a user

represented by the traffic entity”

- Claims 5, 14 – “hidden context represents a task that a user represented by the traffic entity is planning on accomplishing”
- Claims 6, 15 – “hidden context represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity”
- Claims 7, 16 – “hidden context represents a goal of a user represented by the traffic entity, wherein the user expects to achieve the goal within a threshold time interval”
- Claims 9 and 18 – “the sensor data represents one or more images captured by a camera mounted on the autonomous vehicle or a lidar scan captured by a lidar mounted on the autonomous vehicle”

206. Each of the above limitations recites a data element (an image, a parameter, data, etc.) and then defines that data element based on its content.

207. I understand that printed matter is only entitled to patentable weight if it is functionally related to the substrate on which the information is present. In my opinion, the above printed matter limitations do not have any such functional relationship with the “substrate,” which in this case is computer memory. *See* 10[pre] and 19[pre]. Rather than impacting how the computer memory operates, the above limitations merely describe the claimed information.

VII. Conclusion


208. It is my opinion that the Challenged Claims be canceled as unpatentable pursuant to 35 U.S.C. §318(b).

Declaration of Dr. Jason Janét Regarding U.S. Patent No. 11,520,346

* * *

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

Executed on: October 10, 2025



Jason Janét