

UNITED STATES PATENT AND TRADEMARK OFFICE

---

BEFORE THE PATENT TRIAL AND APPEAL BOARD

---

TESLA, INC.

Petitioner

v.

PERCEPTIVE AUTOMATA LLC

Patent Owner

---

IPR2025-01576

U.S. Patent No.: 11,467,579 B2

---

**DECLARATION OF DR. JASON JANÉT  
REGARDING U.S. PATENT NO. 11,467,579**

**TABLE OF CONTENTS**

|             |                                                                                                                                                                                    |           |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b>I.</b>   | <b>INTRODUCTION .....</b>                                                                                                                                                          | <b>1</b>  |
| A.          | Engagement .....                                                                                                                                                                   | 1         |
| B.          | Background and Qualifications .....                                                                                                                                                | 1         |
| C.          | Compensation and Prior Testimony .....                                                                                                                                             | 11        |
| D.          | Information Considered.....                                                                                                                                                        | 11        |
| <b>II.</b>  | <b>LEGAL STANDARDS FOR PATENTABILITY .....</b>                                                                                                                                     | <b>12</b> |
| <b>III.</b> | <b>BACKGROUND TO THE '579 PATENT .....</b>                                                                                                                                         | <b>20</b> |
| A.          | Prosecution History .....                                                                                                                                                          | 20        |
| B.          | Technology Background .....                                                                                                                                                        | 20        |
| C.          | The Alleged Invention.....                                                                                                                                                         | 22        |
| <b>IV.</b>  | <b>LEVEL OF ORDINARY SKILL IN THE ART .....</b>                                                                                                                                    | <b>24</b> |
| <b>V.</b>   | <b>CLAIM CONSTRUCTION.....</b>                                                                                                                                                     | <b>24</b> |
| <b>VI.</b>  | <b>THE CHALLENGED CLAIMS ARE UNPATENTABLE.....</b>                                                                                                                                 | <b>25</b> |
| A.          | Ground 1: Claims 1-20 are unpatentable under 35 U.S.C. over<br>Djuric and AA .....                                                                                                 | 25        |
| 1.          | Summary of Prior Art .....                                                                                                                                                         | 25        |
| 2.          | Motivation to Combine Djuric and AA .....                                                                                                                                          | 29        |
| 3.          | Claim 1 .....                                                                                                                                                                      | 32        |
| 4.          | Claim 2: The method of claim 1, wherein the values describing<br>statistical distribution for each feature comprise a mean value<br>and a standard deviation for the feature.....  | 49        |
| 5.          | Claim 3: The method of claim 1, wherein training the<br>probabilistic neural network comprises determining evidence<br>lower bound between the plurality of values predicted and a |           |

- plurality of values determined from user responses obtained from users presented with images from a training dataset. ....50
6. Claim 4: The method of claim 1, wherein determining the measure of uncertainty for each of the plurality of values comprises: generating a plurality of samples for the input image using the probabilistic neural network; and determining a confidence interval for each of the plurality of values using the plurality of samples.....52
  7. Claim 5: The method of claim 1, wherein navigating the autonomous vehicle comprises ensuring that the autonomous vehicle stays at least a threshold distance away from the traffic entity displayed in the new image, the threshold distance determined based on the measure of uncertainty generated by the probabilistic neural network. ....54
  8. Claim 6: The method of claim 5, wherein the threshold distance is determined to be a value directly related to the measure of uncertainty generated by the probabilistic neural network.....57
  9. Claim 7: The method of claim 1, wherein the hidden context represents a state of mind of a user represented by the traffic entity.....58
  10. Claim 8: The method of claim 1, wherein the hidden context represents a task that a user represented by the traffic entity is planning on accomplishing. ....59
  11. Claim 9: 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.....59
  12. Claim 10: 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. ....60
  13. Claim 11: The method of claim 1, wherein navigating the autonomous vehicle comprises: generating signals for controlling the autonomous vehicle based on motion parameters

|     |                                                                                                                                                                                                                                                                                                                                                                                            |    |
|-----|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
|     | and the hidden context of the traffic entity; and sending the generated signals to controls of the autonomous vehicle. ....                                                                                                                                                                                                                                                                | 62 |
| 14. | Claim 12: The method of claim 1, wherein the probabilistic neural network is a probabilistic convolutional neural network. ....                                                                                                                                                                                                                                                            | 63 |
| 15. | Claim 13 .....                                                                                                                                                                                                                                                                                                                                                                             | 64 |
| 16. | Claim 14: The non-transitory computer readable storage medium of claim 13, wherein the values describing statistical distribution for each feature comprise a mean value and a standard deviation for the feature. ....                                                                                                                                                                    | 66 |
| 17. | Claim 15: The non-transitory computer readable storage medium of claim 13, wherein training the probabilistic neural network comprises determining evidence lower bound between the plurality of values predicted and a plurality of values determined from user responses obtained from users presented with images from a training dataset. ....                                         | 66 |
| 18. | Claim 16: The non-transitory computer readable storage medium of claim 13, wherein determining the measure of uncertainty for each of the plurality of values comprises: generating a plurality of samples for the input image using the probabilistic neural network; and determining a confidence interval for each of the plurality of values using the plurality of samples. ....      | 67 |
| 19. | Claim 17: The non-transitory computer readable storage medium of claim 13, wherein navigating the autonomous vehicle comprises ensuring that the autonomous vehicle stays at least a threshold distance away from the traffic entity displayed in the new image, the threshold distance determined based on the measure of uncertainty generated by the probabilistic neural network. .... | 67 |
| 20. | Claim 18: The non-transitory computer readable storage medium of claim 17, wherein the threshold distance is determined to be a value directly related to the measure of uncertainty generated by the probabilistic neural network. ....                                                                                                                                                   | 67 |

Declaration of Dr. Jason Janét regarding U.S. Patent No. 11,467,579

|     |                                                                                                                                                                                                                                                                                                                                                 |    |
|-----|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 21. | Claim 19: The non-transitory computer readable storage medium of claim 13, wherein navigating the autonomous vehicle comprises: generating signals for controlling the autonomous vehicle based on motion parameters and the hidden context of the traffic entity; and sending the generated signals to controls of the autonomous vehicle..... | 68 |
| 22. | Claim 20.....                                                                                                                                                                                                                                                                                                                                   | 68 |
| B.  | Ground 2: Claims 1-20 are unpatentable under 35 U.S.C. § 103 under Djuric, AA, and Cox .....                                                                                                                                                                                                                                                    | 70 |
| 1.  | Summary of Prior Art: Cox-1008 .....                                                                                                                                                                                                                                                                                                            | 70 |
| 2.  | Motivation to Combine Djuric, AA, and Cox .....                                                                                                                                                                                                                                                                                                 | 72 |
| 3.  | Claim 1 .....                                                                                                                                                                                                                                                                                                                                   | 75 |
| 4.  | Claims 13 and 20.....                                                                                                                                                                                                                                                                                                                           | 80 |
| 5.  | Claims 2-12 and 14-19.....                                                                                                                                                                                                                                                                                                                      | 80 |
| C.  | Grounds 3 and 4: Claims 1-20 are unpatentable under 35 U.S.C. § 103 over Djuric, AA, and Rolfe and/or Djuric, AA, Cox, and Rolfe. ....                                                                                                                                                                                                          | 80 |
| 1.  | Summary of Prior Art: Rolfe-EX1009.....                                                                                                                                                                                                                                                                                                         | 81 |
| 2.  | Motivation to Combine Djuric, AA, Cox, and Rolfe .....                                                                                                                                                                                                                                                                                          | 81 |
| 3.  | Claims 1, 2, 4-14, and 16-20.....                                                                                                                                                                                                                                                                                                               | 82 |
| 4.  | Claim 3: The method of claim 1, wherein training the probabilistic neural network comprises determining evidence lower bound between the plurality of values predicted and a plurality of values determined from user responses obtained from users presented with images from a training dataset. ....                                         | 82 |
| 5.  | Claim 15: The non-transitory computer readable storage medium of claim 13, wherein training the probabilistic neural network comprises determining evidence lower bound between the plurality of values predicted and a plurality of values determined                                                                                          |    |

|             |                                                                                                                                                                              |           |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
|             | from user responses obtained from users presented with images<br>from a training dataset. ....                                                                               | 83        |
| D.          | Grounds 5 and 6: Claims 1-20 are unpatentable under 35 U.S.C. §<br>103 over Djuric, AA, and Ross and/or Djuric, AA, Cox, and Ross ...                                        | 83        |
| 1.          | Summary of Prior Art: Ross-EX1015.....                                                                                                                                       | 83        |
| 2.          | Motivation to Combine Djuric, AA, Cox, and Ross.....                                                                                                                         | 84        |
| 3.          | Claims 1-6, 8, and 10-20.....                                                                                                                                                | 86        |
| 4.          | Claim 7: The method of claim 1, wherein the hidden context<br>represents a state of mind of a user represented by the traffic<br>entity.....                                 | 86        |
| 5.          | Claim 9: The method of claim 1, wherein the hidden context<br>represents a degree of awareness of the autonomous vehicle by a<br>user represented by the traffic entity..... | 89        |
| E.          | Printed Matter.....                                                                                                                                                          | 90        |
| <b>VII.</b> | <b>CONCLUSION .....</b>                                                                                                                                                      | <b>92</b> |

## **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. 11,467,579 (“the ’579 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 1012, 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

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 '579 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 '579 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.

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

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 '579 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.

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.

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. BACKGROUND TO THE '579 PATENT**

#### **A. Prosecution History**

56. I understand that the Office issued a first action notice of allowance for the '579 Patent, stating that the “prior art of record fail[ed] to teach or suggest” the following claim elements:

“generating output representing hidden context for the traffic entity, the output comprising a plurality of values, each value representing a likelihood of receiving a particular user response from a user presented with the image,

...

determining a measure of uncertainty for each of the plurality of values;  
and

navigating the autonomous vehicle to avoid the traffic entity displayed in the new image, the navigation based on at least the measure of uncertainty generated by the probabilistic neural network.”

EX1002, 124-131.

57. But in my opinion, these limitations—like the Challenged Claims as a whole—were well known and should not have been allowed.

#### **B. Technology Background**

58. In my opinion, the '579 Patent generally relates to probabilistic neural networks, which are computer programs that analyze input data and generate an output based on the data. Such models are “probabilistic” because, for example, they generate outputs with a particular statistical distribution. In my opinion, it is

well-known that such distributions have uncertainties associated with them, both in the field of neural networks and more generally in the field of statistics.

59. It is my opinion that neural networks, including probabilistic neural networks, 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’ decisions when labeling training data.

60. For example, it is my opinion that it was commonly known in the art that, for training a traffic prediction probabilistic neural network, human annotators label all instances of training data images that display objects moving in traffic. The labels used by the human annotators correspond to the output predictions made by the neural network. Therefore, after training the neural network with the labels, the neural network can be used to determine the probability of different movement patterns by vehicles represented in new, non-training data, such as camera images. It is my opinion that the neural network’s predictions for vehicles reflect how the human annotators would have labeled the new images because the model was trained based on the human annotators’ decisions when labeling training data. For example, in my opinion, different humans may label images with some degree of variation, which would be reflected by the output from the probabilistic neural network.

**C. The Alleged Invention**

61. It is my opinion that the alleged invention of the '579 Patent relates to (1) training a probabilistic neural network (PNN) to generate output representing “hidden context,” (2) determining a measure of uncertainty of the output, and (3) navigating an autonomous vehicle based on the determined uncertainty. EX1001, 2:6-18, 6:45-63, 7:37-55, 14:11-37, 15:11-26. But, in my opinion, and as illustrated by the cited art, this invention was well-known for many years in the fields of both probabilistic neural networks and autonomous vehicle control and reflects standard training and operations of conventional probabilistic networks. *See* Section VI.

62. In my opinion, the substance of the '579 Patent was well-known. The Patent explains how a PNN is trained to predict “hidden context” for traffic entities, such as vehicles and pedestrians, but this complicated-sounding language is just jargon; such “hidden context” may simply “represent[] the behavior of the traffic entities,” for example. EX1001, Abstract, 2:1-4, 2:64-3:15. The PNN 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 it is my opinion that this concept was well-known. EX1001, 2:10-18, 4:22-35, 4:63-5:3, 7:4-21, 8:39-44, Figs. 4, 8; EX1004, ¶0053; EX1005, ¶0035 (explaining how manual or supervised training techniques may be used to train models to predict future behavior). Indeed, in my opinion, the *inventor's own prior art* (the Cox reference,

EX1008) shows that the '579 Patent should not have issued. Cox discloses to ensure that machine learning model predictions are “more consistent with the decisions of the human annotators” and “better mimic human performance.” EX1008, ¶¶0004, 0032.

63. Thereafter, the trained PNN is used to make a prediction based on “live” input data, such as images of traffic entities. EX1001, 10:15-28, 13:26-48. In my opinion, making predictions based on live input data was well known, for example, at least before 2005. *See generally* EX1011, pp. 250, 353; *see also* EX1004, ¶¶0030-0032, 0061-0068, 0100-0102, 0111-0112; EX1005, ¶¶0073, 0079-0081 (describing the input of current images into a model that predicts future movement of objects depicted in the images.) The Patent describes determining a measure of uncertainty for the PNN output, which was also well known before 2005, in my opinion. *See, e.g.,* EX1011, pp.25-26 (“The key idea in probabilistic robotics is to represent uncertainty explicitly”). EX1001, 5:62-65, 6:45-7:3, 15:27-65, 19:35-48; EX1004, ¶¶0034, 0085-0088, 0093-0094, 0112, 0136.

64. Next, the Patent describes controlling an autonomous vehicle based on the prediction and the measure of uncertainty. EX1001, 2:36-40, 16:48-50, 18:60-65. But it is my opinion that control of “cars that travel by themselves” using the well-known concept of “probabilistic algorithms” was also a well-known concept dating back to at least 2005. *See, e.g.,* EX1011 pp. 24-26, 351, 391, 494, 508. Indeed,

in my opinion, using probabilistic predictions of a neural network to control an autonomous vehicle was well-known, as demonstrated by at least Djuric and AA. EX1004, ¶¶0008, 0021-0024, 0035-0036, 0050, 0078, 0095, 0103-0104, 0114; EX1005, ¶¶0002, 0030, 0050, 0073, 0084.

65. Thus, it is my opinion that the '579 Patent merely describes well-known machine learning concepts. It is my opinion that the Challenged Claims should not have been issued and should now be canceled.

#### **IV. LEVEL OF ORDINARY SKILL IN THE ART**

66. In my opinion, a person having ordinary skill in the art ("POSITA") relevant to the '579 Patent as of February 6, 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**

67. It is my understanding that claim terms in an IPR are construed according to their "ordinary and customary meaning" to those of skill in the art. *Phillips v. AWH Corp.*, 415 F.3d 9 1303 (Fed. Cir. 2005) (en banc); 37 C.F.R.

§42.100(b). It is my understanding that, for purposes of this proceeding, Petitioner does not present any constructions and that Petitioner does not concede that the terms of the challenged claims are definite under 35 U.S.C. §112.

## **VI. THE CHALLENGED CLAIMS ARE UNPATENTABLE**

### **A. Ground 1: Claims 1-20 are unpatentable under 35 U.S.C. over Djuric and AA**

68. In my opinion, Djuric alone renders obvious most limitations of the Challenged Claims, as explained below. Further, it is my opinion that AA explicitly discloses: (1) the description of the neural network as a “probabilistic neural network” (*see* element 1[a]), and (2) the discussion of a “feature vector” in element 1[a-ii], and it would have been obvious to a POSITA to combine the teachings of Djuric and AA. Sections VI.A.3.b, VI.A.3.d (elements 1[a] and 1[a-ii]).

#### **1. Summary of Prior Art**

##### **a. Djuric-EX1004**

69. It is my opinion that 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 that captures camera images of its surroundings, which include objects such as other vehicles and pedestrians, and generates predictions of object motion using a predictive model, such as a neural network. EX1004, ¶¶0027, 0047, 0053, 0061-0067. The model is trained using supervised training techniques with manually

labeled object trajectories. 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," which, in my opinion, demonstrates that the model produces an output closest to the human-labeled trajectories of the training data. EX1004, ¶0094.

70. 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. The model also generates a trajectory confidence level for each of the plurality of predicted trajectories. EX1004, ¶¶0089-0093. The trajectory confidence levels indicate "a per-trajectory probability, quantifying how likely an object is to follow a certain trajectory." EX1004, ¶0089.

71. Additionally, the predicted trajectories each include a plurality of way-points. EX1004, ¶0083. Djuric's neural network model determines "a way-point

confidence level for each way-point,” where the way-point confidence levels for a predicted trajectory “help quantify the confidence in how exactly the object is going to follow a certain trajectory.” EX1004, ¶0085.

72. 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.

73. Thus, it is my opinion that Djuric is analogous art to the ’579 Patent, from the same field of endeavor as the Patent (e.g., predictive neural network model techniques), 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”). EX1004, Abstract, ¶¶0032-0034, 0081; EX1001, 1:38-62, 4:8-22.

**b. AA-EX1005**

74. In my opinion, AA is directed to using “probabilistic generative models in an artificial neural network to predict subject body poses and subject movement intent.” EX1005, ¶0001. AA describes using the movement intent predictions “in autonomous driving applications” to predict pedestrian movement. EX1005, ¶¶0050, 0073.

75. Specifically, AA describes training a convolutional neural network that is a “trained probabilistic model.” EX1005, ¶¶0005-0008, 0034-0035, 0052-0058, 0065-0073. In my opinion, AA’s teachings are consistent with the disclosure of ’579 Patent, which explains that its neural network can be a “probabilistic convolutional neural network.” EX1001, 5:56-60.

76. AA also discloses “feature vectors” that “may include a number that corresponds to a possible feature of the image[.]” EX1005, ¶¶0034-0037. In my opinion, AA further describes that applying a softmax function to one of these feature vectors can convert the numbers in the feature vector “to a probability” for each feature represented the values in the vector. EX1005, ¶0037. Additionally, it is my opinion that AA’s application of a softmax function to the generated feature vector is consistent with the purported invention of the ’579 Patent, which describes processing a feature vector by applying a softmax function to produce probability values. EX1005, 18:23-33, 19:7-10, 19:29-31, Fig. 9, Fig. 10.

77. Thus, it is my opinion that AA is analogous art to the ’579 Patent, from the same field of endeavor as the Patent (e.g., predictive neural network model techniques), 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”). EX1005, ¶¶0034-0044, 0050, Fig. 2D; EX1001, 1:38-62, 4:8-22.

## 2. Motivation to Combine Djuric and AA

78. It is my opinion that a POSITA would have been motivated to combine Djuric and AA such that:

- The model in the Djuric-AA combination teaches a **probabilistic** convolutional neural network, as taught by Djuric and AA. EX1004, ¶¶0027, 0033, 0053, 0069, 0072, 0079-0094, 0100, 0110-0113, 0126, 0132-0139; EX1005, ¶¶0005-0008, 0034-0035, 0052-0058, 0065-0073.
- The model in the Djuric-AA combination generates a feature vector as described in AA. EX1005, ¶¶0034-0037.

79. In my opinion, a POSITA would have been motivated to adopt AA's disclosure of a "probabilistic model" in Djuric for several reasons. For example, Djuric teaches that its model may be a Bayesian network, and AA's disclosures regarding probabilistic model simply provide more implementation detail. EX1004, ¶¶0126, 0132; *see also id.*, 0052-0053, 0081-0094, 0113. Therefore, it is my opinion that a POSITA would have recognized that implementing the probabilistic approach from AA would help Djuric "account for the uncertainty of future movements" and "account for the uncertainty of plausible future poses in predicting subject motion ... in real-time." EX1005, ¶¶0051, 0052, 0065, 0084-0086. Thus, it is my opinion that a POSITA would have been motivated to combine AA's PNN with Djuric's model, which itself describes various examples of PNN's.

80. Djuric even details probabilistic training targets (e.g., Gaussian waypoint uncertainty and per-trajectory classification with “one additional output ... to estimate the per-trajectory probability (resulting in  $5N + 1$  output for a trajectory)”). EX1004, ¶¶0086-0090. Therefore, in my opinion, refining Djuric’s Bayesian network or convolutional neural network as a probabilistic convolutional neural network, per AA is a natural, predictable refinement that aligns with Djuric’s approach. In my opinion, this refinement would also improve Djuric’s prediction results, as AA explains that its probabilistic generative approach “account[s] for the uncertainty of plausible future poses in predicting subject motion.” EX1005, ¶¶0051-0052.

81. In my opinion, a POSITA would also implement AA’s “feature vector” stage to make Djuric’s model more effective. It is my opinion that, by using a compact feature-vector representation after Djuric’s rasterized, multi-channel inputs spatial redundancy would be reduced, thereby improving compute efficiency. This directly serves Djuric’s stated goal that “the combined data set can allow the model’s network to be smaller and, thus, use less computational resources (e.g., processing, storage, etc.).” EX1004, ¶0037. It is my opinion that implementation would be straightforward, because Djuric already “input[s] the combined data set into the machine learned model” to predict trajectories of objects around the autonomous vehicle (EX1004, ¶0037), and it would be straightforward to use this data to generate

a feature vector as taught by AA. Thus, it is my opinion that this modification of Djuric would have been well within the skill of a POSITA, and a POSITA would have an expectation of success in the modification.

82. It is my opinion that, for the foregoing reasons, the combination of Djuric and AA would have been the predictable combination of well-known prior art elements (e.g., Djuric's autonomous vehicle and convolutional neural network; AA's probabilistic convolutional neural network and feature vector generation) according to known methods to yield predictable results (an autonomous vehicle with a predictive convolutional neural network; the Djuric-AA model using a feature vector). In my opinion, the combination would have further used known techniques (known probabilistic convolutional neural networks as taught by AA) to improve similar devices (Djuric's autonomous vehicle and convolutional neural network) in the same way (using AA's aforementioned probabilistic neural network).

83. In my opinion, this combination would also have been well within the skill of a POSITA, and a POSITA would have had a reasonable expectation of success. It is my opinion that a POSITA would have understood that it would be a simple incorporation of the "probabilistic model" of AA in the model of Djuric to achieve the results touted by AA, at least because Djuric discloses that its "machine-learned model can be or can otherwise include one or more various model(s)" to "observe its surrounding environment" and AA discloses a "probabilistic model"

capable of accounting for the uncertainty of plausible future movements of detected objects. EX1004, ¶¶0003, 0027, 0053, 0126, 0132; EX1005, ¶¶0051, 0052, 0065, 0084–0086. It is my opinion that it would be similarly within the skill of a POSITA to incorporate a feature vector as taught by AA, especially if the probabilistic model from AA is being used. Similarly, it is my opinion that a POSITA would have an expectation of success in making the combination with AA because Djuric already discloses that similar models (e.g., a Bayesian network) may be used.

### **3. Claim 1**

84. In my opinion, Djuric and AA render obvious claim 1 for the foregoing reasons.

#### **a. 1[Pre]: A method for navigating autonomous vehicles, comprising:**

85. It is my opinion that Djuric and AA teach the preamble, to the extent it is limiting. For example, 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 for navigating autonomous vehicles*). EX1004, ¶¶0020-0024, 0035-0036, 0044-0046, 0050, 0078, 0095, 0103-0104, 0114, 0118, Fig. 9 (item 920).

**b. 1[a]: training a probabilistic neural network, the probabilistic neural network configured to perform steps comprising:**


86. It is my opinion that Djuric and AA teach element 1[a] because Djuric discloses that its system uses “a model,” which is trained by “supervised training techniques.” EX1004, ¶¶0027, 0053, 0132-0139. Djuric’s trained model may be various types of “neural network[s],” including a “Bayesian network[,],” which is a type of probabilistic neural network (PNN). EX1004, ¶¶0027, 0053, 0126, 0132. The model is configured to perform a number of steps, including those discussed, *infra*, in Sections VI.A.3.c (element 1[a-i]) through VI.A.3.e (element 1[a-iii]) (*the probabilistic neural network configured to perform steps comprising*).

87. To the extent not explicitly disclosed by Djuric, AA describes training a convolutional neural network, as a “trained probabilistic model,” that “generates a probability distribution” to “account for the uncertainty of plausible future poses in predicting subject motion.” EX1005, ¶¶0005-0008, 0034-0035, 0052-0058, 0065-0073. Thus, AA’s convolutional neural network that is trained as a probabilistic model is *a probabilistic neural network*. *Id.* AA’s teachings align with the ’579 Patent, which explains that its neural network can be a “probabilistic convolutional neural network.” EX1001, 5:56-60. In view of this teaching and as described above, it is my opinion that it would have been obvious to a POSITA to implement Djuric using a PNN as taught by AA (*a probabilistic neural network*). As explained above,


Djuric further teaches using “supervised training techniques” on the model (*training a probabilistic neural network*). EX1004, ¶¶0027, 0033, 0053, 0069, 0072, 0079-0094, (“training the model 136 to evaluate which trajectory of a set of trajectories is closest to a ground truth”) 0100, 0110-0113, 0126, 0132-0139; EX1005, ¶¶0005-0008, 0034-0035, 0052-0058, 0065-0073.

**c. 1[a-i]: [the probabilistic neural network configured to perform steps comprising:] receiving as input, an image of traffic, the image displaying a traffic entity belonging to the traffic,**

88. It is my opinion that Djuric and AA teach element 1[a-i] because Djuric discloses that autonomous vehicle 104 includes cameras that generate image data. EX1004, ¶¶0022, 0028-0035, 0046-0059, 0095-0103, 0104-0114, Figs. 1, 2A-2C, 6, 9. The image data includes “objects that are within the surrounding environment of the vehicle 104,” where the objects in the images include “vehicles, pedestrians, [and] bicycles” that are part of traffic in a roadway and on a sidewalk; and the image data, as part of a “combined data set,” is input into the model (*[the probabilistic neural network configured to perform steps comprising:] receiving as input, an image of traffic, the image displaying a traffic entity belonging to the traffic*). *Id.*; *see id.*, ¶¶0021-0023, 0028-0032 (“[t]he prediction system can input the combined data set (e.g., the image data) into the machine-learned model”), 0051, 0059-0063, 0079-0081; Section VI.A.3.b (element 1[a]).



EX1004, Fig. 1 (annotated)



EX1004, Fig. 2A

89. In my opinion, Djuric and AA teach element 1[a-i] because Djuric describes that image data is input (or received) into the model for training purposes as well as during operation. EX1004, ¶¶0021-0022, 0027-0037, 0046-0063, 0079-0094, 0095-0103, 0104-0114, 0134-0139, Figs. 1, 2A-2C, 6, 9; Section VI.A.3.b (element 1[a]). For example, Djuric explains that “sensor data,” which includes image data from cameras, is used during training of the model. EX1004, ¶¶0053, 0079, 0134-0135. Moreover, “supervised training techniques” using, for example, “a set of labeled training data,” and the “backwards propagation of errors” are used to train the model. EX1004, ¶¶0027, 0053, 0134. And further, data from previous

events (e.g., “a previously **observed** trajectory of the object”) are used as training data, which “allow[s] the training data to train a model based on previously recorded real-world object motion and data associated therewith.” *Id.*, ¶0053 (emphasis added); *see id.*, ¶¶0021, 0027, 0053, 0079, 0086-0103, 0134-0135.

- d. **1[a-ii]: [the probabilistic neural network configured to perform steps comprising:] generating a feature vector for a plurality of features, the feature vector comprising values describing statistical distribution for each feature, and**

90. It is my opinion that Djuric and AA teach element 1[a-ii] because AA discloses the model can generate a “first feature vector” and a “second feature vector,” each of which “may include a number that corresponds to a possible feature of the image 226[.]” (*[the probabilistic neural network configured to perform steps comprising:] generating a feature vector for a plurality of features*). EX1005, ¶0034-0037; Section VI.A.3.b (element 1[a]). AA further discloses that a softmax function may convert the numbers of a feature vector to a probability. EX1005, ¶0037. Thus, it is my opinion that, because the numbers of the feature vector can be used to generate a “probability” for each feature, these values “*describ[e] statistical distribution for each feature.*” In my opinion, applying a softmax function to the generated feature vector is consistent with the ’579 Patent, which also describes processing a feature vector by applying a softmax function to produce probability values. EX1001, 18:23-33, 19:7-10, 19:24-33.

91. As described above in claim element 1[a-i], the image data in Djuric includes traffic entities “that are within the surrounding environment of the vehicle 104.” Section VI.A.3.c (element 1[a-i]); EX1004, ¶¶0022, 0028-0035, 0046-0059, 0095-0103, 0104-0114. Similar to AA, Djuric discloses that there are “features associated” with each object (i.e., traffic entity). EX1004, ¶¶0022, 0028-0035, 0057; EX1005, ¶¶0034-0037. Accordingly, it is my opinion that the Djuric–AA combination teaches that each object in the image data may have one or more feature vectors generated (*[the probabilistic neural network configured to perform steps comprising:] generating a feature vector for a plurality of features*) and a softmax function applied to those feature vectors will produce probabilities for each of the features (*the feature vector comprising values describing statistical distribution for each feature*).

- e. **1[a-iii]: [the probabilistic neural network configured to perform steps comprising:] generating output representing hidden context for the traffic entity, the output comprising a plurality of values, each value representing a likelihood of receiving a particular user response from a user presented with the image;**

92. It is my opinion that Djuric and AA teach element 1[a-iii] because Djuric teaches that the model generates an output including both predicted trajectories for an object (*[the probabilistic neural network configured to perform steps comprising:] generating output representing hidden context for the traffic*

*entity*) and associated confidence levels for each trajectory (*the output comprising a plurality of values*). EX1004, ¶¶0005-0007, 0021, 0026-0027, 0032-0043, 0046-0056, 0071-0084, 0085-0095, 0096-0103, 0104-0114, 0126-0139, Claims 1, 2, Figs. 6, 9; Section VI.A.3.b (element 1[a]).

93. Djuric also teaches *generating output representing hidden context for the traffic entity* by teaching that the model generates predicted trajectories in consideration of “the nominal pathway(s) of an object.” EX1004, ¶¶0081-0084. Using such nominal pathways “can allow the model 136 to consider a ‘goal’ that an object may aim to accomplish.” *Id.*, ¶¶0032-0043, 0081-0085, 0107-0113. Djuric further describes “goal-based prediction” where the nominal pathways correspond to “common patterns of object travel.” EX1004, ¶¶0032-0043, 0081-0085, 0107-0113. Djuric’s disclosure of “goal-based prediction” aligns with the ’579 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, 2:64-3:3.

94. The model output in Djuric also *compris[es] a plurality of values* because the output includes “a trajectory confidence level for **each** predicted trajectory.” EX1004, ¶¶0082, 0089-94 (emphasis added). Because Djuric teaches

generating a “plurality of predicted trajectories for an object” there would also be a plurality of confidence levels. EX1004, ¶0082.

95. It is also my opinion that, for several reasons, Djuric’s trajectory confidence levels (*each value*) also “*represent[] a likelihood of receiving a particular user response from a user presented with the image.*” First, the trajectory confidence levels each “indicate a per-trajectory **probability** [(*a likelihood*)], quantifying how likely an object is to follow a certain trajectory.” EX1004, ¶¶0088-0093 (emphasis added). Second, Djuric teaches using “supervised training techniques,” including using “labeled” data. *Id.*, ¶0053; *see also id.*, ¶¶0027, 0053, 0090, 0134-0136. It is my opinion that such labeled data used in supervised training is data labeled by a human, *e.g.*, a “*user response from a user presented with the image.*”


96. It is my opinion that, based on the training techniques disclosed in Djuric, a POSITA would have understood that the model in Djuric is trained such that the predicted probabilities for a given trajectory align with the probability that a human labeler would apply if given the same image. *Id.*; EX1004, ¶¶0021, 0027, 0053, 0079, 0085-0103, 0134-0136. In particular, Djuric teaches that during training, a loss function is “optimized” to “train the network to output plurality of trajectories.” *Id.*, ¶0092. It is my opinion that loss functions were well known in the art as being used determine the error between a model output and the human-labeled training

data, and ultimately used to train the model to more accurately reflect human-labeled training data. EX1004, ¶¶0027, 0053, 0085-0094, 0134-0136. Djuric further discloses that training the model “can include training the model 136 to evaluate which trajectory of a set of trajectories is closest to a ground truth.” EX1004, ¶0094. Therefore, it is my opinion that a POSITA would have understood that “ground truth”—in the context of a supervised trained model, using labeled data—reflects how a human should respond when presented with an image. Thus, in my opinion, a POSITA would have understood that the purpose of this training in Djuric is to bring the model’s predicted trajectories and confidence levels in line with how humans would likely respond to the same image (*each value representing a likelihood of receiving a particular user response from a user presented with the image*).

**f. 1[b]: receiving a new image captured by a camera mounted on an autonomous vehicle navigating through traffic;**

97. It is my opinion that Djuric and AA teach element 1[b] because Djuric discloses that an autonomous vehicle 104 includes cameras that generate image data at various time points, including new image data, as the vehicle navigates through traffic (*receiving a new image captured by a camera mounted on an autonomous vehicle navigating through traffic*). EX1004, ¶¶0021-0023, 0028-0032 (“[t]he

prediction system can input the combined data set (e.g., the image data) into the machine-learned model”), 0046-0059, 0096-0103, 0104-0114, Figs. 1, 2A-2C, 6, 9.



EX1004, Fig. 1 (annotated)

98. Further, it is my opinion that Djuric and AA teach element 1[b], because Djuric’s vehicle computing system 102 processes a plurality of images taken at times  $T_1$ ,  $T_2$ , and  $T_3$ , for example. EX1004, ¶¶0056-0063. The images may be generated with input “image data from a front-facing camera [or] other cameras,” *Id.* The images can be of an object 202 “within a relevant portion of the surrounding environment of the vehicle 102,” including a roadway and lane boundaries. *Id.*, ¶0061, Figs. 2A-2C, 3A-3C.


- g. 1[c]: executing the probabilistic neural network to generate output representing hidden context for at least a traffic entity displayed in the new image;**

99. It is my opinion that Djuric and AA teach element 1[c], because as explained above for element 1[a], Djuric and AA teach training a *probabilistic neural network*. Section VI.A.3.b (element 1[a]). As explained above for element 1[a-iii], Djuric teaches that the model is configured to predict trajectories (*hidden context*) for objects in an image when executed. Section VI.A.3.e (element 1[a-iii]); EX1004, ¶¶0005-0007, 0021, 0026-0027, 0032-0043, 0046-0056, 0071-0084, 0085-0095, 0096-0103, 0104-0114, 0126-0139, Claims 1, 2, Figs. 6, 9. Djuric teaches that the model may be executed over new images captured at times  $T_1$ ,  $T_2$ , and  $T_3$  “to determine a predicted trajectory of an object” in the images. EX1004, ¶¶0070-0071. Thus, Djuric and AA teach element 1[c] in my opinion, because, for the reasons discussed in element 1[a-iii], when new images are provided to the model after the model is trained (as explained with respect to element 1[b]), the trained model will output predicted trajectories for the traffic entity detected in the new images (*executing the probabilistic neural network to generate output representing hidden context for at least a traffic entity displayed in the new image*). *Id.*; Sections VI.A.3.e, VI.A.3.f (elements 1[a-iii], 1[b]).

**h. 1[d]: determining a measure of uncertainty for each of the plurality of values; and**

100. In my opinion, Djuric and AA teach element 1[d] because, as discussed in Sections VI.A.3.e (element 1[a-iii]) and VI.A.3.g (element 1[c]), executing the model generates predicted trajectories and associated trajectory confidence levels for objects (e.g., vehicles, bicycles, or pedestrians) in a new image of the traffic surroundings of autonomous vehicle 104. *Id.*

101. It is my opinion that Djuric describes determining a measure of uncertainty for each of the predicted trajectories and trajectory confidence levels by measuring a “way-point confidence level” each point along the trajectory (*determining a measure of uncertainty for each of the plurality of values*). EX1004, ¶¶0081-0088, 0112. Figure 8 (reproduced below) shows an example of each of the way-point confidence levels, where each way-point confidence level is reflected in red and the predicted trajectories for a vehicle are reflected in green. EX1004, ¶¶0081-0085, Fig. 8.



EX1004, Fig. 8 (annotated)

102. Djuric explains that “[e]ach of the predicted trajectories includes a plurality of predicted waypoints of the object.” EX1004, ¶¶0007, 0083-0094, 0112, 0136, Claims 3, 4, 17, 19, and 20. According to Djuric, “each way-point can be associated with a respective way-point confidence level.” *Id.*, ¶¶0085-0086. These way-point confidence levels “can be indicative of **the per-way-point uncertainty** for the way-point (e.g., predicted future location) within a predicted trajectory” and “can help quantify the confidence” in how exactly the object (the vehicle, pedestrian, bicycle) is going to follow a certain trajectory. *Id.* (parenthetical in original,

emphasis added). Therefore, it is my opinion that, because each way-point confidence level quantifies the confidence in an object following a predicted trajectory, the way-point confidence levels are a determined measure of uncertainty for the predicted trajectory and associated trajectory confidence level (*determining a measure of uncertainty for each of the plurality of values*). *Id.*; Section VI.A.3.e (element 1[a-iii]).

103. It is my opinion that Djuric also teaches element 1[d] in an additional, independent way. Specifically, Djuric calculates a “classification loss” for predicted trajectories, which “quantifies the accuracy of the per-trajectory probability” (i.e., the “trajectory confidence level”). EX1004, ¶¶0090-0091, 0136. Specifically, Djuric’s model is “trained according to the loss function[]” that includes the classification loss of the trajectories, and the “loss function” is the difference between the training data and model outputs. EX1004, ¶¶0090-0091, 0136. Thus, when the model is executed, the classification loss in Djuric is a *measure of uncertainty* for each of the determined predicted trajectories and trajectory confidence levels. EX1004, ¶0090.


- i. **1[e]: navigating the autonomous vehicle to avoid the traffic entity displayed in the new image, the navigation based on at least the measure of**

**uncertainty generated by the probabilistic neural network.**

104. In my opinion, Djuric and AA teach element 1[e] because Djuric explains that “[g]iven knowledge of its surrounding environment, the autonomous vehicle can navigate through such surrounding environment.” EX1004, ¶0003. Djuric’s autonomous vehicle uses a “motion planning system, which can, for example, weigh each of the predicted paths and confidence levels when planning vehicle motion.” EX1004, ¶¶0024, 0034-0035. After the combination’s model produces a plurality of predicted trajectories for an object (e.g., a vehicle, pedestrian, or bicyclist) in image data, including the new image data of Section VI.A.3.f (element 1[b]), vehicle computing system 102 controls autonomous vehicle 104 “based at least in part” on the output predicted trajectories, where the vehicle uses the predicted future locations of the object within a predicted trajectory “to plan and control its motion (e.g., to avoid the object)” (*navigating the autonomous vehicle to avoid the traffic entity displayed in the new image*). EX1004, ¶¶0021, 0034, 0036, 0077, 0083-0094, 0102-0114, Figs. 6, 9 (parenthetical in original); Sections VI.A.3.g (element 1[c]), VI.A.3.h (element 1[d]).

105. As explained for element 1[d], it is my opinion that Djuric’s teaching of each way-point confidence level forms *the measure of uncertainty*. VI.A.3.h (element 1[d]). Djuric teaches that this way-point confidence level is used to predict object trajectories. EX1004, ¶0094 (“[A] predicted trajectory can be selected based

at least in part on one or more of the trajectory confidence levels and/or **the waypoint confidence levels.**”) (emphasis added). These trajectories are then used in the navigation of the autonomous vehicle (e.g., avoiding crossing the traffic entity’s trajectory at the same time, thereby causing a collision), which, in my opinion, means “the navigation [is] based on at least the measure of uncertainty generated by the probabilistic neural network.” *Id.*, ¶¶0083-0094, 0104-0114, Fig. 9.



EX1004, Fig. 8 (annotated)

106. It is my opinion that Djuric also meets element 1[e] in an additional, independent way. As explained for element 1[d], it is my opinion that Djuric's classification loss for trajectory confidence levels also provides a *measure of uncertainty*. See VI.A.3.h (element 1[d]). Djuric's model is trained to minimize the determined classification loss. EX1004, ¶¶0090-0092. This means that the output predicted trajectories in Djuric are based on the determined classification loss (*measure of uncertainty*), and Djuric therefore controls its vehicle to avoid the object based on the determined classification loss (*navigating the autonomous vehicle to avoid the traffic entity displayed in the new image, the navigation based on at least the measure of uncertainty generated by the probabilistic neural network*). EX1004, ¶¶0021, 0036, 0077, 0083-0094, 0104-0114, Fig. 9.

**4. Claim 2: The method of claim 1, wherein the values describing statistical distribution for each feature comprise a mean value and a standard deviation for the feature.**

107. Djuric and AA render obvious claim 2. As discussed for element 1[a-ii], AA teaches a softmax function is applied to the numbers in the disclosed feature vector that “convert[s] the numbers” to a “probability;” the values of the feature vector therefore describe a probability distribution (i.e., probability for and conversely probability against) for each feature represented by each of the values (*the values describing statistical distribution for each feature*). EX1005, ¶0037; Section VI.A.3.d (element 1[a-ii]). It is my opinion that a POSITA would have found

it obvious for such values to *comprise a mean value and a standard deviation for the feature*. This is at least because a mean and a standard deviation are well-known statistical measures that are routinely applied by POSITAs to data. *Id.* Moreover, Djuric teaches that its model can be trained using “an assumption [...] that each way-point can be sampled from a 2D-Gaussian distribution, where the Gaussian distribution is parametrized by 2D mean (e.g., x-y positions) and 2x2 covariance matrix (e.g., modeling uncertainty).” EX1004, ¶0086 (parenthetical in original). In my opinion, in addition to directly disclosing the use of a *mean value*, a POSITA would have understood that the referenced “Gaussian distribution” includes a *standard deviation*. In view of this teaching about the training of the model in Djuric, it is my opinion that the Djuric-AA combination therefore teaches that the values from the feature vector taught by AA would also *comprise a mean value and a standard deviation for the feature*.

**5. Claim 3: The method of claim 1, wherein training the probabilistic neural network comprises determining evidence lower bound between the plurality of values predicted and a plurality of values determined from user responses obtained from users presented with images from a training dataset.**

108. It is my opinion that Djuric and AA render obvious claim 3 because, as discussed for element 1[a-iii], Djuric describes that the model training uses “supervised training techniques” using, for example, “a set of labeled training data,” where the “backwards propagation of errors” is used to train the model. EX1004,

¶¶0027, 0053, 0134-0136; Section VI.A.3.e (element 1[a-iii]). Additionally, Djuric explains that model training is performed by using “a number of sets of data from previous events (e.g., previously obtained state data, previous obtained data associated with a geographic area)” and “a previously observed trajectory of the object,” as well as “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, etc.)” are used as training data, which “allow[s] the training data to train a model based on previously recorded real-world object motion and data associated therewith.” EX1004, ¶0053; *see id.*, ¶¶0021, 0027, 0053, 0079, 0086-0103, 0134-0136. In my opinion, a POSITA would thus have found it obvious that the training data in Djuric would include “*a plurality of values determined from user responses obtained from users presented with images from a training dataset.*”

109. Djuric also discloses model training so that “the negative log-likelihood of the training data set can be minimized.” EX1004, ¶0086; *see also id.*, ¶¶0085-0094, 0136. In my opinion, the use of an Evidence Lower Bound (“ELBO”) was well-known for determining a lower bound of the log-likelihood of a data set to reduce the loss of a model. Thus, based on Djuric’s stated goals, it is my opinion that a POSITA would have understood to implement an ELBO to optimize Djuric’s model (*training the probabilistic neural network comprises determining evidence lower bound between the plurality of values predicted and a plurality of values*

*determined from user responses obtained from users presented with images from a training dataset*). See EX1004, ¶¶0086, 0090-0092.

110. It is my opinion that a POSITA would also have found it obvious to implement an ELBO to optimize Djuric's model based on Djuric's stated goals (*Supra*). EX1004, ¶0086. It is also my opinion that a POSITA would have found it obvious to implement an ELBO to optimize Djuric's model because Djuric discloses that its model may be a Bayesian network, and the determination of an ELBO is an element of variational Bayesian methods. EX1004, ¶¶0126, 0132. Modifying Djuric's model to implement an ELBO would reduce the loss of the model when training, thereby improving the accuracy of the model. *Id.*


**6. Claim 4: The method of claim 1, wherein determining the measure of uncertainty for each of the plurality of values comprises: generating a plurality of samples for the input image using the probabilistic neural network; and determining a confidence interval for each of the plurality of values using the plurality of samples.**

111. It is my opinion that Djuric and AA render obvious claim 4 because, as explained above, Djuric describes determining way-point confidence levels as a measure of uncertainty for each predicted trajectory and the trajectory's associated trajectory confidence level (*determining the measure of uncertainty for each of the plurality of values*). Section VI.A.3.h (element 1[d]).

112. It is my opinion that Djuric and AA render obvious claim 4 because, to generate the way-point confidence levels (*measure of uncertainty*), Djuric

determines the waypoints of predicted trajectories as elements of predicted trajectories that are based on the input image (*generating a plurality of samples for the input image using the probabilistic neural network*). See Sections VI.A.3.b through VI.A.3.e (elements 1[a] through 1[a-iii]).

113. In my opinion, Djuric also teaches that generating the way-point confidence levels includes *determining a confidence interval for each of the plurality of values using the plurality of samples*. Specifically, Djuric teaches that way-point confidence levels can be determined by assuming that each waypoint can be sampled from a 2D-Gaussian distribution “parameterized by [a] 2D mean (e.g., x-y positions) and  $2 \times 2$  covariance matrix (e.g., modeling uncertainty).” EX1004, ¶¶0085-0087 (parentheticals in original). The variance and covariance for each waypoint can also “be represented by a circle centered on [each] way-point 820A, where the radius of the circle is indicative of the confidence.” EX1004, ¶0088, Fig. 8 (below). Therefore, because the variance and covariance for the way-point confidence level describe a range of confidence in the location of an object, it is my opinion that the variance and covariance for each of the waypoints are “*a confidence interval for each of the plurality of values using the plurality of samples*.”



EX1004, Fig. 8 (annotated)


**7. Claim 5: The method of claim 1, wherein navigating the autonomous vehicle comprises ensuring that the autonomous vehicle stays at least a threshold distance away from the traffic entity displayed in the new image, the threshold distance determined based on the measure of uncertainty generated by the probabilistic neural network.**

114. It is my opinion that Djuric and AA render obvious claim 5 because Djuric discloses that its autonomous vehicle uses an object's predicted object trajectories and "one or more predicted future locations of the object" to "plan and control its motion (e.g., to avoid the object)." EX1004, ¶0021 (parenthetical in original); *see also id.*, ¶0036 (listing additional benefits "enhancing

passenger/vehicle safety and improving vehicle efficiency by reducing collisions.”). Djuric further teaches that the autonomous vehicle is navigated “relative to the objects 202, 302 at their **future** locations.” *Id.*, ¶0077 (emphasis added). It is my opinion that, based on Djuric’s stated goal of “reducing collisions,” a POSITA would have understood that, to navigate to avoid a collision with an object, the autonomous vehicle would navigate itself to stay at least a non-zero distance away from the objects in the new image data and their predicted trajectories. (Section VI.A.3.f (element 1[b]), (*wherein navigating the autonomous vehicle comprises ensuring that the autonomous vehicle stays at least a threshold distance away from the traffic entity displayed in the new image*). EX1004, ¶¶0021, 0034, 0036, 0077, 0083-0094, 0104-0114, Fig. 9 (parenthetical in original); 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); Section VI.A.3.i (element 1[e])). Failing to maintain a non-zero minimum distance from a predicted trajectory would result in a collision.

115. Moreover, as discussed in Section VI.A.3.h (element 1[d]), Djuric explains that “[e]ach of the predicted trajectories includes a plurality of predicted waypoints of the object.” EX1004, ¶¶0007, 0083-0094, 0112, 0136, Claims 3, 4, 17, 19, and 20. A “way-point confidence level” (the *measure of uncertainty*) is

determined “for each way-point within a predicted trajectory” such that “each way-point can be associated with a respective way-point confidence level.” *Id.*, ¶¶0085-0086. These way-point confidence levels can be represented by a circle with a radius that is “indicative of the confidence/uncertainty associated with that particular predicted future location of the object.” *Id.*, ¶0088, Fig. 8 (annotated below with way-point confidence in red).



EX1004, Fig. 8 (annotated)

116. Because the autonomous vehicle in Djuric is navigated based on the predicted trajectories along with the way-point confidence levels (*see* Section

VI.A.3.i (element 1[e]); EX1004, ¶0094) it is my opinion that it would have been obvious to a POSITA, based on the teachings of Djuric, that the non-zero distance that Djuric's autonomous vehicle would have stayed away from objects in the new image data would have been based on the way-point confidence levels (*the threshold distance determined based on the measure of uncertainty generated by the probabilistic neural network*). EX1004, ¶¶0007, 0077, 0083-0094, 0112, 0136, Fig. 8, Claims 3, 4, 17, 19, and 20.

**8. Claim 6: The method of claim 5, wherein the threshold distance is determined to be a value directly related to the measure of uncertainty generated by the probabilistic neural network.**

117. It is my opinion that Djuric and AA render obvious claim 6 because, as discussed for claim 5, the non-zero distance Djuric's autonomous vehicle is kept away from the objects in image data, including the new image data (Section VI.A.3.f (element 1[b])), and their predicted trajectories form *the threshold distance*, and the way-point confidence levels form *the measure of uncertainty generated by the Djuric-AA-Cox probabilistic neural network*. Section VI.A.3.h (element 1[d]); Section VI.A.7 (claim 5). It is also my opinion that it would have been obvious to a POSITA that—consistent with basic driving techniques—the autonomous vehicle should give an object more space when there is more uncertainty about the object's predicted location. Thus, in my opinion, the threshold distance maintained from an object should increase when the measure of uncertainty increases (*wherein the*

*threshold distance is determined to be a value directly related to the measure of uncertainty*). Section VI.A.3.h (element 1[d]); Section VI.A.7 (claim 5).

**9. Claim 7: The method of claim 1, wherein the hidden context represents a state of mind of a user represented by the traffic entity.**

118. It is my opinion that Djuric and AA render obvious claim 7 because, as explained for claim 1, Djuric teaches that the predicted trajectories and associated trajectory confidence levels for an object—for example, for a vehicle—represent a *hidden context* for the vehicle because they represent goals the vehicle “may aim to accomplish” that are associated with “nominal pathways,” such as turning, merging, or continuing straight. EX1004, ¶¶0032-0043, 0081-0085, 0107-0113; Section VI.A.3.g (element 1[c]); Section VI.A.3.e (element 1[a-iii]). It is my opinion that, 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. In my opinion, this is consistent with the ’579 Patent’s explanation that a state of mind may include “that a pedestrian is likely to cross the street.” EX1001, 12:11-16.

**10. Claim 8: The method of claim 1, wherein the hidden context represents a task that a user represented by the traffic entity is planning on accomplishing.**

119. It is my opinion that Djuric and AA render obvious claim 8 because, as explained for claim 1, Djuric teaches that the predicted trajectories and associated trajectory confidence levels for an object—for example, for a vehicle—represent a *hidden context* for the vehicle because they represent goals the vehicle “may aim to accomplish” that are associated with “nominal pathways,” such as turning, merging, or continuing straight. EX1004, ¶¶0032-0043, 0081-0085, 0107-0113; Section VI.A.3.g (element 1[c]); Section VI.A.3.e (element 1[a-iii]). It is my opinion that, 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. In my opinion, this is consistent with the ’579 Patent’s explanation that a task that a user is planning to accomplish may be “cross[ing] the street.” EX1001, 3:3-15; 12:11-16.

**11. Claim 9: 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.**

120. It is my opinion that Djuric and AA render obvious claim 9 because, as explained for claim 1, Djuric teaches that the predicted trajectories and associated trajectory confidence levels for an object—for example, for a vehicle—represent a

*hidden context* for the vehicle because they represent goals the vehicle “may aim to accomplish” that are associated with “nominal pathways,” such as turning, merging, or continuing straight. EX1004, ¶¶0032-0043, 0081-0085, 0107-0113; Section VI.A.3.g (element 1[c]); Section VI.A.3.e (element 1[a-iii]). It is my opinion that, because this *hidden context* represents the goals/intentions of a vehicle (and in turn, the vehicle’s driver), 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. In my opinion, a POSITA would have understood that another vehicle’s (e.g., object 202 in Figure 5) awareness of the autonomous vehicle 104 would impact “the potential risk to the vehicle 104” that is relevant to vehicle 104 passing object 202. *See* EX1004, ¶0076. It is my opinion that this is also consistent 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. *Id.*; EX1004, ¶0077.

**12. Claim 10: 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.**

121. It is my opinion that Djuric and AA render obvious claim 10, because as explained for claim 1, Djuric teaches that the predicted trajectories and associated trajectory confidence levels for an object—for example, for a vehicle—represent a

*hidden context* for the vehicle because they represent goals the vehicle “may aim to accomplish” that are associated with “nominal pathways,” such as turning, merging, or continuing straight. EX1004, ¶¶0032-0043, 0081-0085, 0107-0113; Section VI.A.3.g (element 1[c]); Section VI.A.3.e (element 1[a-iii]). It is my opinion that, because this *hidden context* represents the goals/intentions of a vehicle (and in turn, the vehicle’s driver), 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.

122. Moreover, Djuric discloses each trajectory “can include a plurality of predicted future locations of the object 302 over time (e.g., at future times),” which teaches the *goal* is expected to be achieved within a *threshold time interval* because the future times form an interval 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 would have been obvious to a POSITA that the limited computational resources define the time horizon of Djuric’s predictions. Any prediction Djuric makes of a goal that a user will achieve occurs within the computational window or time horizon (*threshold time*) defined by Djuric’s limited computational resources. Thus, Djuric teaches *the user expects to achieve the goal within a threshold time interval*.

**13. Claim 11: The method of claim 1, wherein navigating the autonomous vehicle comprises: generating signals for controlling the autonomous vehicle based on motion parameters and the hidden context of the traffic entity; and sending the generated signals to controls of the autonomous vehicle.**

123. It is my opinion that Djuric and AA render obvious claim 11, because, as described in claim 1, Djuric describes that the autonomous vehicle is navigated (*navigating the autonomous vehicle*). Section VI.A.3.i (element 1[e]). 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).

124. It is my opinion that Djuric further teaches that the autonomous vehicle is controlled *based on motion parameters and the hidden context of the traffic entity* for several reasons. First, Djuric teaches that “the vehicle computing system 102 can control the motion of the vehicle 104 based at least in part on the output 406 from the model 136.” EX1004, ¶0078; *see also id.* ¶¶0024, 0034-0035. Second, Djuric teaches that “[t]he output 406 can be indicative of a predicted trajectory of the object” (*hidden context*) and “the output 406 can include a predicted velocity of the object” (*motion parameters*). EX1004, ¶¶0071, 0102.

125. In my opinion, Djuric discloses an additional *motion parameter* 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. 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. Section VI.A.3.g (element 1[c]); EX1004, ¶¶0104-0114. Thus, in my opinion, the navigation of the autonomous vehicle includes controlling the autonomous vehicle based on the predicted trajectories and also based on the state data describing object motion (*controlling the autonomous vehicle based on motion parameters and the hidden context of the traffic entity*). EX1004, ¶¶0021-0025, 0034-0036, 0051, 0071, 0077-0078, 0083-0094, 0102-0114, 0127, Fig. 9; VI.A.3.i (element 1[e]).

**14. Claim 12: The method of claim 1, wherein the probabilistic neural network is a probabilistic convolutional neural network.**

126. It is my opinion that Djuric and AA render obvious claim 12 because, as explained for element 1[a], AA teaches that the model may be a probabilistic neural network. Section VI.A.3.b (element 1[a]); EX1005, ¶¶0005-0008, 0034-0035,

0052-0058, 0065-0073. Djuric further teaches that “[n]eural networks can include convolutional neural networks.” EX1004, ¶0027; *see also id.*, ¶¶0033, 0053, 0069, 0072, 0079-0094, 0100, 0110-0113, 0126, 0132-0139. Therefore, it is my opinion that the Djuric-AA combination teaches that *the probabilistic neural network is a probabilistic convolutional neural network. Id.*

### 15. Claim 13

127. It is my opinion that Djuric and AA render obvious claim 13 for the same reasons discussed with respect to claim 1, as indicated below. *See* Section VI.A.3 (claim 1).

**a. 13[Pre]: A non-transitory computer readable storage medium storing instructions, that when executed by a processor, cause the processor to perform steps comprising:**

128. *See* Section VI.A.3.a (element 1[Pre]). Additionally, it is my opinion that Djuric teaches element 13[Pre] because 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. 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 steps comprising*). *Id.*, ¶0117-0118.

- b. 13[a]: training a probabilistic neural network, the probabilistic neural network configured to perform steps comprising:**

129. *See* Section VI.A.3.b (element 1[a]).

- c. 13[a-i]: receiving as input, an image of traffic, the image displaying a traffic entity belonging to the traffic,**

130. *See* Section VI.A.3.c (element 1[a-i]).

- d. 13[a-ii]: generating a feature vector for a plurality of features, the feature vector comprising values describing statistical distribution for each feature, and**

131. *See* Section VI.A.3.d (element 1[a-ii]).

- e. 13[a-iii]: generating output representing hidden context for the traffic entity, the output comprising a plurality of values, each value representing a likelihood of receiving a particular user response from a user presented with the image;**

132. *See* Section VI.A.3.e (element 1[a-iii]).

- f. 13[b]: receiving a new image captured by a camera mounted on an autonomous vehicle navigating through traffic;**

133. *See* Section VI.A.3.f (element 1[b]).

- g. 13[c]: executing the probabilistic neural network to generate output representing hidden context**

**for at least a traffic entity displayed in the new image;**

134. *See* Section VI.A.3.g (element 1[c]).

**h. 13[d]: determining a measure of uncertainty for each of the plurality of values; and**

135. *See* Section VI.A.3.h (element 1[d]).

**i. 13[e]: navigating the autonomous vehicle to avoid the traffic entity displayed in the new image, the navigation based on at least the measure of uncertainty generated by the probabilistic neural network.**

136. *See* Section VI.A.3.i (element 1[e]).

**16. Claim 14: The non-transitory computer readable storage medium of claim 13, wherein the values describing statistical distribution for each feature comprise a mean value and a standard deviation for the feature.**

137. It is my opinion that Djuric renders obvious claim 14 for the same reasons discussed above in Section VI.A.4 (claim 2).

**17. Claim 15: The non-transitory computer readable storage medium of claim 13, wherein training the probabilistic neural network comprises determining evidence lower bound between the plurality of values predicted and a plurality of values determined from user responses obtained from users presented with images from a training dataset.**

138. It is my opinion that Djuric renders obvious claim 15 for the same reasons discussed above in Section VI.A.5 (claim 3).

**18. Claim 16: The non-transitory computer readable storage medium of claim 13, wherein determining the measure of uncertainty for each of the plurality of values comprises: generating a plurality of samples for the input image using the probabilistic neural network; and determining a confidence interval for each of the plurality of values using the plurality of samples.**

139. It is my opinion that Djuric renders obvious claim 16 for the same reasons discussed above in Section VI.A.6 (claim 4).

**19. Claim 17: The non-transitory computer readable storage medium of claim 13, wherein navigating the autonomous vehicle comprises ensuring that the autonomous vehicle stays at least a threshold distance away from the traffic entity displayed in the new image, the threshold distance determined based on the measure of uncertainty generated by the probabilistic neural network.**

140. It is my opinion that Djuric renders obvious claim 17 for the same reasons discussed above in Section VI.A.7 (claim 5).

**20. Claim 18: The non-transitory computer readable storage medium of claim 17, wherein the threshold distance is determined to be a value directly related to the measure of uncertainty generated by the probabilistic neural network.**

141. It is my opinion that Djuric renders obvious claim 18 for the same reasons discussed above in Section VI.A.8 (claim 6).

**21. Claim 19: The non-transitory computer readable storage medium of claim 13, wherein navigating the autonomous vehicle comprises: generating signals for controlling the autonomous vehicle based on motion parameters and the hidden context of the traffic entity; and sending the generated signals to controls of the autonomous vehicle.**

142. Djuric renders obvious claim 19 for the same reasons discussed above in Section VI.A.13 (claim 11).

**22. Claim 20**

143. It is my opinion that Djuric and AA render obvious claim 20 for the same reasons presented above with respect to claims 1 and 13. *See* Sections VI.A.3, VI.A.15 (claims 1, 13).

**a. 20[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:**

144. *See* Section VI.A.15.a (element 13[Pre]). Djuric discloses the “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)” (*A computer system*). EX1004, ¶0116; *see also id.*, ¶¶0040, 0045, 0140.

**b. 20[a]: training a probabilistic neural network, the probabilistic neural network configured to perform steps comprising:**

145. *See* Section VI.A.3.b (element 1[a]).

- c. **20[a-i]: [the probabilistic neural network configured to perform steps comprising:] receiving as input, an image of traffic, the image displaying a traffic entity belonging to the traffic,**

146. *See* Section VI.A.3.c (element 1[a-i]).

- d. **20[a-ii]: [the probabilistic neural network configured to perform steps comprising:] generating a feature vector for a plurality of features, the feature vector comprising values describing statistical distribution for each feature, and**

147. *See* Section VI.A.3.d (element 1[a-ii]).

- e. **20[a-iii]: [the probabilistic neural network configured to perform steps comprising:] generating output representing hidden context for the traffic entity, the output comprising a plurality of values, each value representing a likelihood of receiving a particular user response from a user presented with the image;**

148. *See* Section VI.A.3.e (element 1[a-iii]).

- f. **20[b]: receiving a new image captured by a camera mounted on an autonomous vehicle navigating through traffic;**

149. *See* Section VI.A.3.f (element 1[b]).

- g. **20[c]: executing the probabilistic neural network to generate output representing hidden context for at least a traffic entity displayed in the new image;**

150. *See* Section VI.A.3.g (element 1[c]).

- h. 20[d]: determining a measure of uncertainty for each of the plurality of values; and**

151. See Section VI.A.3.h (element 1[d]).

- i. 20[e]: navigating the autonomous vehicle to avoid the traffic entity displayed in the new image, the navigation based on at least the measure of uncertainty generated by the probabilistic neural network.**

152. See Section VI.A.3.i (element 1[e]).

**B. Ground 2: Claims 1-20 are unpatentable under 35 U.S.C. § 103 under Djuric, AA, and Cox**

153. To the extent Djuric alone does not render obvious the use of human labeled data to generate values *representing a likelihood of receiving a particular user response from a user presented with the image* as recited in element 1[a-iii], the claimed operations are also disclosed by Cox. For the reasons described below, it is my opinion that it would have been obvious to a POSITA to combine these features from Cox into Djuric, along with the above described teachings from AA.

**1. Summary of Prior Art: Cox-1008**

154. 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 '579 Patent. EX1001, Cover; EX1008, Cover. In my opinion, Cox is a material reference to the '579 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.

155. Cox is directed toward “machine learning augmented by measurements acquired from human populations.” EX1008, Abstract, ¶¶0001-0004. Cox teaches crowd-sourcing training data used to train machine learning algorithms. EX1008, Abstract, ¶¶0001-0004, 0015-0016, 0025-0034, 0044, 0060. Just like the ’579 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 EX1008, ¶¶0029-0032, Figs. 1A-B, 3 *with* EX1001, 4:23-35, 8:39-44; Fig. 4. The goal of Cox’s machine learning techniques is to “better mimic human performance.” EX1008, ¶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. EX1008, ¶¶ 0056, 0060.

156. In my opinion, Cox is analogous art to the ’579 Patent, from the same field of endeavor as the ’579 Patent (e.g., predictive neural network model techniques), and is reasonably pertinent to the particular problem that the ’579 Patent was trying to solve (e.g., how to “accurately predict motion of non-stationary traffic

objects”), describing the gathering training data from a population of number of human labelers, and training a model using the training data, and executing it to output predictions. EX1008, ¶¶0004, 0029-0032, Figs. 1A-B, 3; EX1001, 1:38-62, 4:8-22.

## **2. Motivation to Combine Djuric, AA, and Cox**

157. It is my opinion that a POSITA would have been motivated to further combine Djuric and AA with Cox such that confidence levels (*each value*) from Djuric, quantifying how likely an object is to follow a certain trajectory (EX1004, ¶¶0088-0093), are generated in accordance with human annotated/labeled training data. As taught by Cox, using human annotated image data to train the Djuric-AA 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.” EX1008, ¶¶0004, 0027-0032, 0051. In my opinion, a POSITA would have been motivated to combine Djuric and AA with Cox in this way for several reasons.

158. First, in my opinion, 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.a. 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. EX1008, ¶¶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.” EX1008, ¶¶0004, 0060; *see also id.*, ¶¶0008-0010, 0015-0019, 0025-0035, 0051. Thus, it is my opinion that Cox provides a natural expansion and improvement on Djuric’s base disclosure.

159. It is my opinion that a POSITA seeking to implement the Djuric-AA system 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. EX1008, ¶¶0016, 0056. Therefore, in my opinion, Cox provides an explicit teaching, suggestion, or motivation to make a combination with a system for controlling driverless (or autonomous) vehicles, just like Djuric.

160. In my opinion, a POSITA would have also been motivated to combine the teachings of the Djuric-AA system and Cox because the combination would have benefitted the Djuric-AA system. It is my opinion that enhancing the Djuric-AA system 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 the Djuric-AA model so it better considers the state of mind of other road users and what they might be thinking when present in autonomous vehicle's environment. EX1008, ¶¶0007-0010, 0015-0019, 0025-0035, 0051, 0060. It is my opinion that this would have improved the Djuric-AA system's goal-based predictions, and thus helped reduce collisions between the autonomous vehicle and human drivers or pedestrians.

161. For all the reasons explained above, it is my opinion that a POSITA would have been motivated to combine Djuric, AA, and Cox as described. In my opinion, this combination would have been the predictable combination of well-known prior art elements (e.g., the Djuric-AA model; Cox's model training and optimization techniques) according to known methods to yield predictable results (the Djuric-AA model being a supervised learning model; the Djuric-AA-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). In my opinion, 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-AA model for goal-based predictions) in the same way (using Cox's aforementioned techniques). Lastly, in my opinion, due to the overlap in the disclosures of Djuric, AA, and Cox and their compatibility, using human annotated training data from Cox with the teachings of Djuric and AA was well within the skill of a POSITA, and a POSITA would have had an expectation of success in this combination.

### 3. Claim 1

162. In my opinion, Cox's teachings only impact how the model generates values *representing a likelihood of receiving a particular user response from a user presented with the image* in element 1[a-iii]. Thus, it is my opinion that the combination of Djuric, AA, and Cox teaches element 1[a-iii] for the reasons explained below, as well as the reasons addressed in Ground 1. *Supra*, Sections VI.A.3, VI.A.3.e (element 1[a-iii]). Moreover, Djuric, AA, and Cox teach the remaining elements of claim 1 for the same reasons addressed in Ground 1. *Id.*

- a. **1[a-iii]: [the probabilistic neural network configured to perform steps comprising:] generating output representing hidden context for the traffic entity, the output comprising a plurality of values, each value representing a**

**likelihood of receiving a particular user response  
from a user presented with the image;**

163. In my opinion, Djuric, AA, and Cox teach 1[a-iii] for the same reasons addressed in Ground 1 (Section VI.A.3.e), and in an additional way. Specifically, and as described further below, Cox teaches that human-labeled image data is used to train a model, so that the model's output *represent[]s a likelihood of receiving a particular user response from a user presented with the image*. EX1008, ¶¶0019, 0029-0038, 0060; Fig. 4B; Section VI.A.3.e (Ground 1, element 1[a-iii]). So, implementing these teachings of Cox with Djuric, the trajectory confidence levels (*each value*) in Djuric would further “*represent[] a likelihood of receiving a particular user response from a user presented with the image.*”

164. Cox uses “[a] crowd-sourc[ed][]" and “large population” of human annotators to annotate a “training object” such as an image, generate “statistical information” based on all of the responses received from the multiple different human annotators, and train a supervised machine learning model using the statistical information. EX1008, ¶¶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.*, ¶0017, Figs. 1A, 1B. For example, the human annotator responses may describe a state of mind of the training object, 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.” EX1008, ¶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.

165. In my opinion, Cox additionally teaches that the model’s output represents *a likelihood of receiving a particular user response from a user* 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.” EX1008, ¶¶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.” EX1008, ¶¶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.” EX1008, ¶¶0032, 0034.

166. Similarly, it is also my opinion that Cox describes that its model is trained to minimize the difference between the model's predicted output and the human-labeled training data. EX1008, ¶¶0019, 0029-0038, 0060; Fig. 4B. Specifically, Cox 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. EX1008, ¶¶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 the decisions of the human annotators." EX1008, ¶¶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. EX1008, ¶¶0019, 0029-0038, 0060; Fig. 4B.

167. Thus, as applied to Djuric and AA, it is my opinion that Cox teaches that the trajectory confidence levels quantifying how likely an object is to follow a certain trajectory (EX1004, ¶¶0088-0093) are generated in accordance with the human-annotated/labeled training data. EX1008, ¶¶0004, 0027-0032, 0051. The human-annotated/labeled training data includes annotated/labeled image data and is used to train the Djuric-AA model to provide predictions that are “more consistent with the decisions of the human annotators,” and that “better mimic human performance.” *Id.* Thus, in my opinion, the model’s outputs are representative of how humans would have annotated or labeled images input into the model if the images were part of the training data. EX1008, ¶¶0004, 0019, 0029-0038, 0051, 0060; Fig. 4B. As such, the Djuric-AA-Cox combination teaches that the model generates predicted trajectories for an object and associated trajectory confidence levels for the predicted trajectories (*[the probabilistic neural network configured to perform steps comprising:] generating output representing hidden context for the traffic entity, the output comprising a plurality of values*) (see Section VI.A.3.e (element 1[a-iii]), and Cox, applied to Djuric and AA, further teaches the predicted confidence levels (*each value*) represent likelihoods of how humans would have annotated/labeled the input image data, thus teaching *each value representing a likelihood of receiving a particular user response from a user presented with the*

*image*. EX1004, ¶¶0021, 0027, 0053, 0079, 0085-0103, 0134-0136; EX1008, ¶¶0004, 0019, 0029-0038, 0051, 0060; Fig. 4B; Section VI.A.3.e (element 1[a-iii]).

#### **4. Claims 13 and 20**

168. It is my opinion that Djuric, AA, and Cox render obvious claims 13 and 20 for the same reasons discussed above with respect to claim 1. The combination with Cox only impacts the analysis for elements 13[a-iii] and 20[a-iii]. Elements 13[a-iii] and 20[a-iii] are taught by Djuric, AA, and Cox for the reasons stated above for 1[a-iii]. *Supra*, Section VI.B.3.a (element 1[a-iii]). Djuric, AA, and Cox further teach elements 13[a-iii] and 20[a-iii], as well as the remaining elements of claims 13 and 20, for the same reasons addressed for Djuric and AA in Ground 1. Sections VI.A.15 (claim 13), VI.A.22 (claim 20).

#### **5. Claims 2-12 and 14-19**

169. It is my opinion that Djuric, AA, and Cox render obvious claims 2-12 and 14-19 for the same reasons discussed with respect to Djuric and AA in Section VI.A above (Ground 1).

#### **C. Grounds 3 and 4: Claims 1-20 are unpatentable under 35 U.S.C. § 103 over Djuric, AA, and Rolfe and/or Djuric, AA, Cox, and Rolfe.**

170. It is my opinion that, to the extent Djuric alone does not render obvious dependent claims 3 and 15, Rolfe confirms that the recited use of an evidence lower bound (ELBO) is a well-known statistical technique. The analysis for independent claims 1 and 13 is unchanged from Grounds 1 and 2. Claims 3 and 15 are thus

rendered obvious over the further combination of Rolfe, with either Djuric and AA (Ground 3), or Djuric, AA, and Cox (Ground 4).

### **1. Summary of Prior Art: Rolfe-EX1009**

171. Rolfe discloses quantum processors that may be used in machine learning, for example, for “Bayesian statistics, and artificial neural networks”. EX1009, Abstract, ¶0004. In this context, Rolfe discusses use of “the evidence lower bound (ELBO)” as an example approach “to optimize a lower bound on the log-likelihood function.” *Id.*, ¶0086.

172. Rolfe is analogous art to the ’579 Patent, from the same field of endeavor as the Patent (e.g., predictive neural network model techniques), 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”). EX1009, Abstract, ¶¶0002-0004, 0086; EX1001, 1:38-62, 4:8-22.

### **2. Motivation to Combine Djuric, AA, Cox, and Rolfe**

173. In my opinion, a POSITA would have been motivated to combine Rolfe with Djuric (along with AA and optionally Cox) to use the disclosed ELBO techniques to minimize the negative log-likelihood of the training data set in Djuric. Djuric teaches that the model is trained so that “the negative log-likelihood of the training data set can be minimized.” EX1004, ¶0086; *see also id.*, ¶¶0085-0094, 0136. Thus, it is my opinion that the disclosure of ELBO in Rolfe simply provides

implementation details (already in the context of machine learning) that a POSITA could utilize to achieve the stated goal in Djuric. EX1009, ¶0086. As such, in my opinion, it would be well within the skill of a POSITA to use the implementation details from Rolfe for ELBO in the context of Djuric, and a POSITA would have an expectation of success in making this combination.

**3. Claims 1, 2, 4-14, and 16-20**

174. It is my opinion that Djuric, AA, and Rolfe (Ground 3) and Djuric, AA, Cox, and Rolfe (Ground 4) render obvious claims 1, 2, 4-14, and 16-20 for the same reasons discussed above with respect to Djuric and AA (Ground 1, *Supra* Section VI.A) and Djuric, AA, and Cox (Ground 2, *Supra* Section VI.B).

**4. Claim 3: The method of claim 1, wherein training the probabilistic neural network comprises determining evidence lower bound between the plurality of values predicted and a plurality of values determined from user responses obtained from users presented with images from a training dataset.**

175. In my opinion, combining Rolfe with Djuric, AA (and potentially Cox) makes express what a POSITA would have already understood from Djuric alone: the use of an Evidence Lower Bound was well-known for determining a lower bound of the log-likelihood of a data set to reduce the loss of a model during training (*wherein training the probabilistic neural network comprises determining evidence lower bound between the plurality of values predicted and a plurality of values determined from user responses obtained from users presented with images from a*

*training dataset*). EX1009, ¶0086. In my opinion, the further combination of Rolfe thus renders obvious Claim 3 for the same reasons discussed in Ground 1, with the added confirmation from Rolfe that a POSITA would already be familiar with the applications of an ELBO to the values recited in Claim 3. *Supra* Section VI.A.5 (claim 3).

**5. Claim 15: The non-transitory computer readable storage medium of claim 13, wherein training the probabilistic neural network comprises determining evidence lower bound between the plurality of values predicted and a plurality of values determined from user responses obtained from users presented with images from a training dataset.**

176. It is my opinion that Djuric, AA, and Rolfe (Ground 3) and Djuric, AA, Cox and Rolfe (Ground 4) render obvious claim 15 for the same reasons discussed above with respect to claim 3. *Supra* Section VI.C.4.

**D. Grounds 5 and 6: Claims 1-20 are unpatentable under 35 U.S.C. § 103 over Djuric, AA, and Ross and/or Djuric, AA, Cox, and Ross**

**1. Summary of Prior Art: Ross-EX1015**

177. Ross is directed to techniques for “maneuvering” an autonomous vehicle. EX1015, 1:30-31, 3:39-42, 3:64-5:5. Ross describes 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) so the autonomous vehicle can “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 it is safe to cross. *Id.*, 4:40-44.

178. Ross is analogous art to the ’579 Patent, from the same field of endeavor as the Patent (e.g., predictive neural network model techniques), 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”); Ross describes an model that provides “[r]easoning about the intents of objects, such as other road users.” *Id.*, 5:6-11, 15:1-16:52.

## **2. Motivation to Combine Djuric, AA, Cox, and Ross**

179. In my opinion, POSITA would have been motivated to combine Djuric, AA, and/or Cox with Ross as discussed in Sections VI.A.2 and VI.B.2. A POSITA would have been further motivated to combine Ross with Djuric and AA (and optionally Cox) such that the Djuric-AA-Ross probabilistic neural network model generates “goal-based prediction[s]” using the example goals disclosed by Djuric as well as the example intents or actions disclosed by Ross.

180. As discussed regarding elements 1[a-iii] and 1[c], Djuric describes generating “goal-based prediction[s]” of trajectories of objects in a vehicle’s environment. EX1004, ¶¶0081-0085; Sections VI.A.3.e, VI.A.3.g (elements 1[a-iii] and 1[c]). 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 "[e]ach possible intent of the set includes a corresponding next action to be initiated by the road user." *Id.* 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.

181. For the foregoing reasons, a POSITA would have been motivated to combine Ross with Djuric and AA (and optionally Cox) as described. This combination would have been the predictable combination of well-known prior art elements (e.g., the Djuric-AA 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 the Djuric-AA 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 (the Djuric-AA 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 an expectation of success in the modification.

**3. Claims 1-6, 8, and 10-20**


182. It is my opinion that Djuric, AA, and Ross (Ground 5) and Djuric, AA, Cox, and Ross (Ground 6) render obvious claims 1-6, 8, and 10-20 for the same reasons discussed above with respect to Djuric and AA (Ground 1, Section VI.A) and Djuric, AA, and Cox (Ground 2, Section VI.B).

**4. Claim 7: The method of claim 1, wherein the hidden context represents a state of mind of a user represented by the traffic entity.**

183. It is my opinion that Djuric, AA, and Ross render obvious claim 7 for the same reasons discussed in Section VI.A.9 (Ground 1, claim 7). To the extent Djuric-AA does not explicitly disclose the hidden context represents a state of mind, the combination of Djuric-AA and Ross renders this obvious.

184. 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)

185. Ross’s teachings align with the ’579 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, 12:11-16; see also EX1015, 16:5-41, Fig. 9.

**5. Claim 9: 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.**

186. Djuric and Ross render obvious claim 9 for the same reasons discussed in Section VI.A.11 (Ground 1, claim 9). To the extent Djuric-AA does not explicitly disclose the hidden context represents a degree of awareness of the autonomous vehicle by a user, the combination of Djuric-AA and Ross renders this obvious.

187. 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)

188. 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 “to cross the road when clear[.]” *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 pedestrian of the vehicle driving toward the pedestrian (*represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity*).

#### **E. Printed Matter**

189. 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.

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

- 1[a-i], 13[a-i], 20[a-i] – “an image of traffic, the image displaying a traffic entity belonging to the traffic”

- 1[a-ii], 13[a-ii], 20[a-ii] – “the feature vector comprising values describing statistical distribution for each feature”
- 1[a-iii], 13[a-iii], 20[a-iii] – “output representing hidden context for the traffic entity, the output comprising a plurality of values, each value representing a likelihood of receiving a particular user response from a user presented with the image”
- 1[c], 13[c], 20[c] – “output representing hidden context for at least a traffic entity displayed in the new image”
- 2, 14 – “values describing statistical distribution for each feature comprise a mean value and a standard deviation for the feature”
- 3, 15 – “user responses obtained from users presented with images from a training dataset”
- 7 – “the hidden context represents a state of mind of a user represented by the traffic entity”
- 8 – “the hidden context represents a task that a user represented by the traffic entity is planning on accomplishing”
- 9 – “the hidden context represents a degree of awareness of the autonomous vehicle by a user represented by the traffic entity”

- 10 – “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”

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

192. 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* 13[pre] and 20[pre]. Rather than impacting how the computer memory operates, the above limitations merely describe the claimed information.

## **VII. Conclusion**

193. 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. 11,467,579

\* \* \*

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 9, 2025



Jason Janét