

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

PERCEPTIVE AUTOMATA, LLC,	)	
	)	
Plaintiff,	)	
v.	)	Civil Action No. 2:25-cv-742-JRG
	)	
TESLA, INC.,	)	
	)	
Defendant.	)	
	)	

DEFENDANT’S NOTICE OF SUBPOENA TO AURORA INNOVATION, INC.

Pursuant to Rule 45 of the Federal Rules of Civil Procedure, Tesla, Inc. (“Tesla”), by and through its counsel, requests documents and a deposition from Aurora Innovation, Inc. pursuant to the attached Subpoena. A true and correct copy of the Subpoena is attached as Attachment A.

Dated: December 10, 2025

Respectfully submitted,

By: /s/ Roger Fulghum
Roger Fulghum
Lead Attorney
Texas State Bar No. 00790724
BAKER BOTTS L.L.P.
910 Louisiana St
Houston, TX 77002
Telephone: 713-229-1707
Facsimile: 713-229-2707
Email:
roger.fulghum@bakerbotts.com

Mark Speegle
Texas State Bar No. 24117198
BAKER BOTTS L.L.P.
401 South First Street, Suite 1300
Austin, Texas 78704
Telephone: 512-322-2536
Facsimile: 512-322-3636
Email:
mark.speegle@bakerbotts.com

Helen Trac
California Bar No. 285824
Tesla, Inc.
1501 Page Mill, Building 6
Palo Alto, CA 94304
Telephone: 628-218-1429
Email: htrac@tesla.com

Ashraf Fawzy
DC Bar No. 989132
Telephone: (202) 905-9221
Email: afawzy@tesla.com
Paul Margulies
DC Bar No. 1000297
Telephone: (202) 695-5388
Email: pmargulies@tesla.com
Tesla, Inc.
800 Connecticut Ave. NW
Washington, DC 20006

Counsel for Tesla, Inc.

CERTIFICATE OF SERVICE

I hereby certify that all counsel of record are being served this 10th day of December 2025, by e-mail at the addresses indicated below:

Andrea Leigh Fair
Miller Fair Henry PLLC
1507 Bill Owens Parkway
Longview, TX 75604
andrea@millerfairhenry.com

Charles Austin Ginnings
Taryn Trusty
Timothy E Grochocinski
Nelson Bumgardner Conroy PC
745 McClintock Road, Suite 340
Burr Ridge, IL 60527
austin@nelbum.com
taryn@nelbum.com
tim@nelbum.com

Janson Westmoreland
John Paul Murphy
Nelson Bumgardner Conroy PC
3131 West 7th Street, Suite 300
Fort Worth, TX 76107
janson@nelbum.com
murphy@nelbum.com

Jonathan Hart Rastegar
Patrick Joseph Conroy
Nelson Bumgardner Conroy PC - Dallas
2727 N. Harwood Street, Suite 250
Dallas, TX 75201
JON@nelbum.com
pat@nelbum.com

/s/ Roger Fulghum
Roger Fulghum

UNITED STATES DISTRICT COURT

for the

Eastern District of Texas

PERCEPTIVE AUTOMATA LLC

Plaintiff

v.
TESLA, INC.

Defendant

Civil Action No. 2:25-cv-00742-JRG

SUBPOENA TO TESTIFY AT A DEPOSITION IN A CIVIL ACTION

To:

Aurora Inovations LLC, 1654 Smallman Street, Pittsburgh, PA 15222

(Name of person to whom this subpoena is directed)

☒ **Testimony:** YOU ARE COMMANDED to appear at the time, date, and place set forth below to testify at a deposition to be taken in this civil action. If you are an organization, you must promptly confer in good faith with the party serving this subpoena about the following matters, or those set forth in an attachment, and you must designate one or more officers, directors, or managing agents, or designate other persons who consent to testify on your behalf about these matters:

Place: Veritext's Regus Location: 2403 Sidney St., Suite 240,
Pittsburgh, PA 15203

Date and Time:
01/12/2026 9:00 am

The deposition will be recorded by this method: Stenographic and videotaped

☒ **Production:** You, or your representatives, must also bring with you to the deposition the following documents, electronically stored information, or objects, and must permit inspection, copying, testing, or sampling of the material:

See Attachment A

The following provisions of Fed. R. Civ. P. 45 are attached – Rule 45(c), relating to the place of compliance; Rule 45(d), relating to your protection as a person subject to a subpoena; and Rule 45(e) and (g), relating to your duty to respond to this subpoena and the potential consequences of not doing so.

Date: 12/10/2025

CLERK OF COURT

OR

/s/ Roger Fulghum

Signature of Clerk or Deputy Clerk

Attorney's signature

The name, address, e-mail address, and telephone number of the attorney representing *(name of party)* Defendant, Tesla, Inc., who issues or requests this subpoena, are:

Roger Fulghum, Baker Botts L.L.P., 910 Louisiana, Houston, TX 77002; 713-229-1707; roger.fulghum@bakerbotts.com

Notice to the person who issues or requests this subpoena

If this subpoena commands the production of documents, electronically stored information, or tangible things before trial, a notice and a copy of the subpoena must be served on each party in this case before it is served on the person to whom it is directed. Fed. R. Civ. P. 45(a)(4).

Perceptive Exhibit 2042
IPR2025-01576

Civil Action No. 2:25-cv-00742-JRG

PROOF OF SERVICE

(This section should not be filed with the court unless required by Fed. R. Civ. P. 45.)

I received this subpoena for *(name of individual and title, if any)* _____
on *(date)* _____.

☐ I served the subpoena by delivering a copy to the named individual as follows: _____

_____ on *(date)* _____; or

☐ I returned the subpoena unexecuted because: _____
_____.

Unless the subpoena was issued on behalf of the United States, or one of its officers or agents, I have also
tendered to the witness the fees for one day's attendance, and the mileage allowed by law, in the amount of
\$ _____.

My fees are \$ _____ for travel and \$ _____ for services, for a total of \$ 0.00 .

I declare under penalty of perjury that this information is true.

Date: _____

Server's signature

Printed name and title

Server's address

Additional information regarding attempted service, etc.:

Federal Rule of Civil Procedure 45 (c), (d), (e), and (g) (Effective 12/1/13)

(c) Place of Compliance.

(1) For a Trial, Hearing, or Deposition. A subpoena may command a person to attend a trial, hearing, or deposition only as follows:

- (A) within 100 miles of where the person resides, is employed, or regularly transacts business in person; or
- (B) within the state where the person resides, is employed, or regularly transacts business in person, if the person
 - (i) is a party or a party's officer; or
 - (ii) is commanded to attend a trial and would not incur substantial expense.

(2) For Other Discovery. A subpoena may command:

- (A) production of documents, electronically stored information, or tangible things at a place within 100 miles of where the person resides, is employed, or regularly transacts business in person; and
- (B) inspection of premises at the premises to be inspected.

(d) Protecting a Person Subject to a Subpoena; Enforcement.

(1) Avoiding Undue Burden or Expense; Sanctions. A party or attorney responsible for issuing and serving a subpoena must take reasonable steps to avoid imposing undue burden or expense on a person subject to the subpoena. The court for the district where compliance is required must enforce this duty and impose an appropriate sanction—which may include lost earnings and reasonable attorney's fees—on a party or attorney who fails to comply.

(2) Command to Produce Materials or Permit Inspection.

(A) *Appearance Not Required.* A person commanded to produce documents, electronically stored information, or tangible things, or to permit the inspection of premises, need not appear in person at the place of production or inspection unless also commanded to appear for a deposition, hearing, or trial.

(B) *Objections.* A person commanded to produce documents or tangible things or to permit inspection may serve on the party or attorney designated in the subpoena a written objection to inspecting, copying, testing, or sampling any or all of the materials or to inspecting the premises—or to producing electronically stored information in the form or forms requested. The objection must be served before the earlier of the time specified for compliance or 14 days after the subpoena is served. If an objection is made, the following rules apply:

- (i) At any time, on notice to the commanded person, the serving party may move the court for the district where compliance is required for an order compelling production or inspection.
- (ii) These acts may be required only as directed in the order, and the order must protect a person who is neither a party nor a party's officer from significant expense resulting from compliance.

(3) Quashing or Modifying a Subpoena.

(A) *When Required.* On timely motion, the court for the district where compliance is required must quash or modify a subpoena that:

- (i) fails to allow a reasonable time to comply;
- (ii) requires a person to comply beyond the geographical limits specified in Rule 45(c);
- (iii) requires disclosure of privileged or other protected matter, if no exception or waiver applies; or
- (iv) subjects a person to undue burden.

(B) *When Permitted.* To protect a person subject to or affected by a subpoena, the court for the district where compliance is required may, on motion, quash or modify the subpoena if it requires:

(i) disclosing a trade secret or other confidential research, development, or commercial information; or

(ii) disclosing an unretained expert's opinion or information that does not describe specific occurrences in dispute and results from the expert's study that was not requested by a party.

(C) *Specifying Conditions as an Alternative.* In the circumstances described in Rule 45(d)(3)(B), the court may, instead of quashing or modifying a subpoena, order appearance or production under specified conditions if the serving party:

- (i) shows a substantial need for the testimony or material that cannot be otherwise met without undue hardship; and
- (ii) ensures that the subpoenaed person will be reasonably compensated.

(e) Duties in Responding to a Subpoena.

(1) Producing Documents or Electronically Stored Information. These procedures apply to producing documents or electronically stored information:

(A) *Documents.* A person responding to a subpoena to produce documents must produce them as they are kept in the ordinary course of business or must organize and label them to correspond to the categories in the demand.

(B) *Form for Producing Electronically Stored Information Not Specified.* If a subpoena does not specify a form for producing electronically stored information, the person responding must produce it in a form or forms in which it is ordinarily maintained or in a reasonably usable form or forms.

(C) *Electronically Stored Information Produced in Only One Form.* The person responding need not produce the same electronically stored information in more than one form.

(D) *Inaccessible Electronically Stored Information.* The person responding need not provide discovery of electronically stored information from sources that the person identifies as not reasonably accessible because of undue burden or cost. On motion to compel discovery or for a protective order, the person responding must show that the information is not reasonably accessible because of undue burden or cost. If that showing is made, the court may nonetheless order discovery from such sources if the requesting party shows good cause, considering the limitations of Rule 26(b)(2)(C). The court may specify conditions for the discovery.

(2) Claiming Privilege or Protection.

(A) *Information Withheld.* A person withholding subpoenaed information under a claim that it is privileged or subject to protection as trial-preparation material must:

- (i) expressly make the claim; and
- (ii) describe the nature of the withheld documents, communications, or tangible things in a manner that, without revealing information itself privileged or protected, will enable the parties to assess the claim.

(B) *Information Produced.* If information produced in response to a subpoena is subject to a claim of privilege or of protection as trial-preparation material, the person making the claim may notify any party that received the information of the claim and the basis for it. After being notified, a party must promptly return, sequester, or destroy the specified information and any copies it has; must not use or disclose the information until the claim is resolved; must take reasonable steps to retrieve the information if the party disclosed it before being notified; and may promptly present the information under seal to the court for the district where compliance is required for a determination of the claim. The person who produced the information must preserve the information until the claim is resolved.

(g) Contempt.

The court for the district where compliance is required—and also, after a motion is transferred, the issuing court—may hold in contempt a person who, having been served, fails without adequate excuse to obey the subpoena or an order related to it.

Attachment A

ATTACHMENT A

Pursuant to Rule 45 of the Federal Rules of Civil Procedure, Tesla is seeking the production of the documents requested below for inspection and copying. The due date for production set in this subpoena is January 12, 2026. Tesla is also seeking deposition testimony from Aurora on the documents and topics set out below. The time and date for the deposition is January 12, 2026 at 9:00 am. Tesla is willing to negotiate with Aurora an alternative date and time for both the production of documents, and a deposition. Materials produced in this litigation will be governed by the Protective Order entered in this case, dated November 21, 2025. *See* Attachment B.

DEFINITIONS

1. The terms “You”, “Your”, and “Aurora” mean Aurora Innovation, Inc., and each of that entity’s parents, subsidiaries, divisions, affiliates, predecessors, assigns, successors, and acquired assets or business units, and any of its present or former officers, directors, trustees, employees, agents, representatives, attorneys, patent agents, and all other persons acting on its behalf.

2. The term “Uber” means Uber Technologies, Inc. and each of that entity’s parents, subsidiaries, divisions, affiliates, predecessors, assigns, successors, and acquired assets or business units, and any of its present or former officers, directors, trustees, employees, agents, representatives, attorneys, patent agents, and all other persons acting on its behalf.

3. The term “Uber Prior Art Vehicles” means any vehicle with self-driving or autonomous features (including without limitation vehicles that implement the motion prediction model described in Attachment C, or similar technology) that was designed, developed, tested, operated, offered for sale, used in commerce, or otherwise used by You or Uber, whether in the

United States or another country before February 6, 2019. *See* Attachment C at 8 (“[T]he framework was subsequently deployed to a fleet of autonomous vehicles.”).

INSTRUCTIONS

1. Tesla is seeking documents and information related to prior art for certain patents that were originally issued to Perceptive Automata, Inc. and later assigned to the plaintiff in this lawsuit, Perceptive Automata, LLC. Tesla is willing to meet and confer and answer any questions concerning the material sought in this subpoena. Please reach out to Roger Fulghum (roger.fulghum@bakerbotts.com or mobile:281-217-5031) or Mark Speegle (mark.speegle@bakerbotts.com or mobile 254-749-9120) with any questions concerning this subpoena.

REQUESTS FOR PRODUCTION

1. Documents sufficient to identify all Uber Prior Art Vehicles that were in public use, publicly demonstrated, offered for sale, used in commerce, or otherwise available to the public before each of the following dates: February 6, 2019; January 30, 2019; and July 5, 2017.
2. For each Uber Prior Art Vehicle, documents sufficient to establish when such Uber Prior Art Vehicle was first in public use, publicly demonstrated, offered for sale, used in commerce, or otherwise available to the public.
3. Documents sufficient to show the operation of each Uber Prior Art Vehicle.
4. Documents sufficient to show the operation of (including inputs and outputs to) all machine learning models, algorithms, routines, methods, functions, or structures in relation to or utilized in each Uber Prior Art Vehicle, including but not limited to the MobileNet-v2 predictor. *See, e.g.*, Attachment C at 6-8 (“Lastly, taking the best performing MNv2 and as base and adding uncertainty outputs led to the best results.”).
5. Documents sufficient to show the sensors (e.g., cameras, radar, and LiDAR) implemented on Uber Prior Art Vehicles.
6. Documents sufficient to show how You or Uber collected, stored, labeled, and/or generated training data used to train machine learning models implemented in or executed by the Uber Prior Art Vehicles. *See, e.g.*, Attachment C at 6.
7. Documents sufficient to show how motion prediction models were “deployed to a fleet of autonomous vehicles.” Attachment C at 8.
8. Documents sufficient to show how “[a]ll models were trained end-to-end from scratch.” Attachment C at 6.

9. Documents sufficient to show how the “model with uncertainty outputs [] was initialized with a corresponding model without uncertainty and then fine-tuned.” Attachment C at 6.

10. Documents sufficient to identify all vehicles, products, or services related to U.S. Patent No. 11,112,796, including all such vehicles, products, or services implementing U.S. Patent No. 11,112,796, whether in whole or in part.

DEPOSITION TOPICS

1. All documents produced in response to this subpoena, including the authenticity thereof and whether each such documents is a record of an act, event, condition, opinion, or diagnosis that (i) was made at or near the time by — or from information transmitted by — someone with knowledge; (ii) was kept in the course of a regularly conducted activity of a business, organization, occupation, or calling, whether or not for profit; and (iii) was made as a regular practice of that activity.

2. The identity of all Uber Prior Art Vehicles that were in public use, publicly demonstrated, offered for sale, used in commerce, or otherwise available to the public before each of the following dates: February 6, 2019; January 30, 2019; and July 5, 2017.

3. The operation of the Uber Prior Art Vehicles, including that of any components or software thereof, in relation to or within any computers, servers, cloud instances, or the like.

4. The operation of (including inputs and outputs to) all machine learning models, algorithms, routines, methods, functions, or structures in relation to or within each Uber Prior Art Vehicle.

5. Attachment C, including the contents and the circumstances around its authorship and publication.

6. The operation and implementation of MobileNet-v2 in relation to the Uber Prior Art Vehicles. *See* Attachment C at 6-8.
7. The sensors (e.g., cameras, radar, and LiDAR) implemented on the Uber Prior Art Vehicles.
8. The collection, storage, labeling, and/or generating of training data used to train machine learning models executed by or implemented in the Uber Prior Art Vehicles.
9. How motion prediction models were “deployed to a fleet of autonomous vehicles.” Attachment C at 8.
10. How “[a]ll models were trained end-to-end from scratch.” Attachment C at 6.
11. How the “model with uncertainty outputs [] was initialized with a corresponding model without uncertainty and then fine-tuned.” Attachment C at 6.
12. All vehicles, products, or services related to U.S. Patent No. 11,112,796, including all such vehicles, products, or services implementing U.S. Patent No. 11,112,796, whether in whole or in part.

Attachment B

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

PERCEPTIVE AUTOMATA LLC,	§	
	§	
<i>Plaintiff,</i>	§	
	§	Case No. 2:25-cv-00742-JRG
v.	§	
	§	JURY TRIAL DEMANDED
TESLA, INC.	§	
	§	
<i>Defendant.</i>	§	

PROTECTIVE ORDER

WHEREAS, Plaintiff Perceptive Automata LLC and Defendant Tesla, Inc., hereafter referred to as “the Parties,” believe that certain information that is or will be encompassed by discovery demands by the Parties involves the production or disclosure of trade secrets, confidential business information, or other proprietary information;

WHEREAS, the Parties seek a protective order limiting disclosure thereof in accordance with Federal Rule of Civil Procedure 26(c):

THEREFORE, it is hereby stipulated among the Parties and ORDERED that:

1. Each Party may designate as confidential for protection under this Order, in whole or in part, any document, information or material that constitutes or includes, in whole or in part, confidential or proprietary information or trade secrets of the Party or a Third Party to whom the Party reasonably believes it owes an obligation of confidentiality with respect to such document, information or material (“Protected Material”). Protected Material shall be designated by the Party producing it by affixing a legend or stamp on such document, information or material as follows: “CONFIDENTIAL.” The word “CONFIDENTIAL” shall be placed clearly on each page of the Protected Material (except deposition and hearing

transcripts) for which such protection is sought. For deposition and hearing transcripts, the word “CONFIDENTIAL” shall be placed on the cover page of the transcript (if not already present on the cover page of the transcript when received from the court reporter) by each attorney receiving a copy of the transcript after that attorney receives notice of the designation of some or all of that transcript as “CONFIDENTIAL.”

2. Any document produced under Patent Rules 2-2, 3-2, and/or 3-4 before issuance of this Order with the designation “Confidential” or “Confidential - Outside Attorneys’ Eyes Only” shall receive the same treatment as if designated “RESTRICTED - ATTORNEYS’ EYES ONLY” under this Order, unless and until such document is redesignated to have a different classification under this Order.
3. With respect to documents, information or material designated “CONFIDENTIAL,” “RESTRICTED - ATTORNEYS’ EYES ONLY,” or “RESTRICTED CONFIDENTIAL - SOURCE CODE” (“DESIGNATED MATERIAL”),¹ subject to the provisions herein and unless otherwise stated, this Order governs, without limitation: (a) all documents, electronically stored information, and/or things as defined by the Federal Rules of Civil Procedure; (b) all pretrial, hearing or deposition testimony, or documents marked as exhibits or for identification in depositions and hearings; (c) pretrial pleadings, exhibits to pleadings and other court filings; (d) affidavits; and (e) stipulations. All copies, reproductions, extracts, digests and complete or partial summaries prepared from any DESIGNATED MATERIALS shall also be considered DESIGNATED MATERIAL and treated as such

¹ The term DESIGNATED MATERIAL is used throughout this Protective Order to refer to the class of materials designated as “CONFIDENTIAL,” “RESTRICTED - ATTORNEYS’ EYES ONLY,” or “RESTRICTED CONFIDENTIAL - SOURCE CODE,” both individually and collectively.

under this Order.

4. A designation of Protected Material (i.e., “CONFIDENTIAL,” “RESTRICTED - ATTORNEYS’ EYES ONLY,” or “RESTRICTED CONFIDENTIAL SOURCE CODE”) may be made at any time. Production of documents, information, or material that has not been designated as DESIGNATED MATERIAL, whether inadvertently or otherwise, shall not be deemed a waiver in whole or in part of a claim for confidential treatment. Any party that inadvertently or otherwise produces Protected Material without designating it as DESIGNATED MATERIAL may request destruction of that Protected Material by notifying the recipient(s), as soon as reasonably possible after the Producing Party becomes aware of the inadvertent or unintentional disclosure, and providing replacement Protected Material that is properly designated. The recipient(s) shall then destroy all copies of the inadvertently or unintentionally produced Protected Materials and any documents, information or material derived from or based thereon.
5. “CONFIDENTIAL” documents, information and material may be disclosed only to the following persons, except upon receipt of the prior written consent of the designating party, upon order of the Court, or as set forth in paragraph 12 herein:
 - (a) Outside Counsel of Record in this Action for the Parties;
 - (b) employees of such counsel assigned to and reasonably necessary to assist such counsel in the litigation of this Action;
 - (c) in-house counsel for the Parties who either have responsibility for making decisions dealing directly with the litigation, or who are assisting outside counsel in the litigation;
 - (d) up to and including three (3) designated representatives of each of the Parties to the extent reasonably necessary for the litigation of this Action, except that either party may in good faith request the other party’s consent to designate one or more additional representatives, the other party shall not unreasonably withhold such consent, and the requesting party may seek leave of Court to designate such

additional representative(s) if the requesting party believes the other party has unreasonably withheld such consent;

- (e) outside consultants or experts (*i.e.*, not existing employees or affiliates of a Party or an affiliate of a Party) retained for the purpose of this litigation, provided that: (1) such consultants or experts are not presently employed by the Parties hereto for purposes other than this Action; (2) before access is given, the consultant or expert has completed the Undertaking attached as Exhibit A hereto and the same is served upon the Producing Party with a current curriculum vitae of the consultant or expert, including a list of other cases from the prior five (5) years in which the individual has provided a report or testified (at trial or at deposition, at least ten (10) days before access to the Protected Material is to be given to that consultant or Undertaking to object to and notify the Receiving Party in writing that it objects to disclosure of Protected Material to the consultant or expert. The Parties agree to promptly confer and use good faith to resolve any such objection. If the Parties are unable to resolve any objection, the objecting Party may file a motion with the Court within fifteen (15) days of the notice, or within such other time as the Parties may agree, seeking a protective order with respect to the proposed disclosure. The objecting Party shall have the burden of proving the need for a protective order. No disclosure shall occur until all such objections are resolved by agreement or Court order;
 - (f) independent litigation support services, including persons working for or as court reporters, graphics or design services, jury or trial consulting services, and photocopy, document imaging, and database services retained by counsel and reasonably necessary to assist counsel with the litigation of this Action; and
 - (g) the Court and its personnel.
6. A Party shall designate documents, information or material as “CONFIDENTIAL” only upon a good faith belief that the documents, information or material contains confidential or proprietary information or trade secrets of the Party or a Third Party to whom the Party reasonably believes it owes an obligation of confidentiality with respect to such documents, information or material.
7. Documents, information or material produced pursuant to any discovery request in this Action, including but not limited to Protected Material designated as DESIGNATED MATERIAL, shall be used by the Parties only in the litigation of this Action and shall not be used for any other purpose. Any person or entity who obtains access to DESIGNATED

MATERIAL or the contents thereof pursuant to this Order shall not make any copies, duplicates, extracts, summaries or descriptions of such DESIGNATED MATERIAL or any portion thereof except as may be reasonably necessary in the litigation of this Action. Any such copies, duplicates, extracts, summaries or descriptions shall be classified DESIGNATED MATERIALS and subject to all of the terms and conditions of this Order.

8. To the extent a Producing Party believes that certain Protected Material qualifying to be designated CONFIDENTIAL is so sensitive that its dissemination deserves even further limitation, the Producing Party may designate such Protected Material “RESTRICTED - ATTORNEYS’ EYES ONLY,” or to the extent such Protected Material includes extremely sensitive DESIGNATED MATERIAL representing computer source code and/or live data (that is, data as it exists residing in a database or databases) as well as associated “Readme” files that are stored in the same directory as the source code in the ordinary course of business (collectively, “Source Code Material”), the Producing Party may designate such Protected Material as “RESTRICTED CONFIDENTIAL - SOURCE CODE.”
9. For Protected Material designated RESTRICTED - ATTORNEYS’ EYES ONLY, access to, and disclosure of, such Protected Material shall be limited to individuals listed in paragraphs 5(a-c) and (e-g); provided, however, that access by in-house counsel pursuant to paragraph 5(c) be limited to in-house counsel who exercise no competitive decision-making authority on behalf of the client.
10. For Protected Material designated RESTRICTED CONFIDENTIAL - SOURCE CODE, the following additional restrictions apply:
 - (a) Access to a Party’s Source Code Material shall be provided only on “stand-alone” computer(s) (that is, the computer may not be linked to any network, including a local area network (“LAN”), an intranet or the Internet) and on which all access ports have been disabled, as necessary and appropriate to prevent and protect

against any unauthorized copying, transmission, removal, or other transfer of any Source Code outside or away from the computer on which the Source Code is provided for inspection. Additionally, except as provided in paragraph 10(u) below, the stand-alone computer(s) may only be located at or near the offices of the Producing Party's Outside Counsel of Record or In-House Counsel of Record if the Party is not represented by outside counsel. The Producing Party may visually monitor the activities of the Receiving Party's representatives during any source code review, but only to ensure that there is no unauthorized recording, copying, or transmission of the source code;

- (b) No recordable media or recordable devices, including without limitation sound recorders, computers, cellular telephones, peripheral equipment, cameras, CDs, DVDs, or drives of any kind shall be permitted into the source code review room.
- (c) Parties shall ensure that appropriate software sufficient for viewing, and searching the Source Code Material is installed on the Secure Source Code Computers. The Receiving Party's counsel and/or experts may request that commercially available software for viewing, and searching source code be installed on the Secure Source Code Computer. The Receiving Party must provide the Producing Party with electronic media containing such software, or links to a public location from which such software may be downloaded, and any necessary licenses at least four (4) days in advance of the date upon which the Receiving Party wishes to have the software tools available for use on the Secure Source Code Computers. Examples of software may include SciTools Understand, Eclipse, Android Studio, Microsoft Visual Studio, Notepad++, UltraEdit, SlickEdit, Beyond Compare, CodeSuite from Software Analysis, PowerGrep, Mozilla Firefox, Adobe Acrobat, or other similar software that provides for viewing, and searching source code;
- (d) The Receiving Party shall make reasonable efforts to restrict its requests for such access to the stand-alone computer(s) to normal business hours, which for purposes of this paragraph shall be 9:00 a.m. through 6:00 p.m. The Parties agree to cooperate in good faith such that maintaining the Producing Party's Source Code Material at or near the offices of its Outside Counsel of Record, or In-House Counsel of Record if the Party is not represented by outside counsel, shall not unreasonably hinder the Receiving Party's ability to efficiently and effectively conduct the prosecution or defense of this Action;
- (e) The Producing Party shall provide the Receiving Party with information explaining how to start, log on to, and operate the stand-alone computer(s) in order to access the produced Source Code Material on the stand-alone computer(s);
- (f) The Producing Party will produce Source Code Material in computer searchable format on the stand-alone computer(s) as described above;
- (g) Access to Protected Material designated RESTRICTED CONFIDENTIAL -

SOURCE CODE shall be limited to Outside Counsel of Record and up to three (3) outside consultants or experts² (i.e., not existing employees or affiliates of a Party or an affiliate of a Party) retained for the purpose of this litigation and approved to access such Protected Materials pursuant to paragraph 5(e) above. A receiving Party may include excerpts of Source Code Material in a pleading, exhibit, expert report, discovery document, deposition transcript, other Court document, provided that the Source Code Documents are appropriately marked under this Order, restricted to those who are entitled to have access to them as specified herein, and, if filed with the Court, filed under seal in accordance with the Court's rules, procedures and orders;

- (h) To the extent portions of Source Code Material are quoted in a Source Code Exhibit, either (1) the entire Source Code Exhibit will be stamped and treated as RESTRICTED CONFIDENTIAL - SOURCE CODE or (2) those pages containing quoted Source Code Material will be separately stamped and treated as RESTRICTED CONFIDENTIAL - SOURCE CODE and, if filed with the Court, filed under seal in accordance with the Court's rules, procedures, and orders;
- (i) Except as set forth in paragraph 10(l) below, no electronic copies of Source Code Material shall be made without prior written consent of the producing Party, except as necessary to create documents which, pursuant to the Court's rules, procedures and order, must be filed or served electronically.
- (j) The Receiving Party shall not be permitted to make its own printouts and photocopies of Source Code Material. At the Receiving Party's request, the Producing Party shall provide up to three (3) copies of all printouts of Source Code Material, all of which shall be designated and clearly labeled "RESTRICTED CONFIDENTIAL SOURCE CODE";
- (k) If the Receiving Party's Outside Counsel of Record, consultants, or experts obtain printouts or photocopies of Source Code Material, the Receiving Party shall ensure that such outside counsel, consultants, or experts keep the printouts or photocopies in a secured locked area in the offices of such outside counsel, consultants, or expert. The Receiving Party may also temporarily keep the printouts or photocopies at: (i) the Court for any proceedings(s) relating to the Source Code Material, for the dates associated with the proceeding(s); (ii) the sites where any deposition(s) relating to the Source Code Material are taken, for the dates associated with the deposition(s); and (iii) any intermediate location reasonably necessary to transport the printouts or photocopies (e.g., a hotel prior to a Court proceeding or deposition); and
- (l) A Producing Party's Source Code Material may only be transported by the

² For the purposes of this paragraph, an outside consultant or expert is defined to include the outside consultant's or expert's direct reports and other support personnel, such that the disclosure to a consultant or expert who employs others within his or her firm to help in his or her analysis shall count as a disclosure to a single consultant or expert.

Receiving Party at the direction of a person authorized under paragraph 10(e) above to another person authorized under paragraph 10(e) above, on paper or removable encrypted electronic media (*e.g.*, a DVD, CD-ROM, or flash memory “stick”) via hand carry, Federal Express or other similarly reliable courier. Source Code Material may not be transported or transmitted electronically over a network of any kind, including a LAN, an intranet, or the Internet. Source Code Material may only be transported electronically for the purpose of Court proceeding(s) as set forth in paragraph 10(k) above and is at all times subject to the transport restrictions set forth herein. But, for those purposes only, the Source Code Materials may be loaded onto a stand-alone computer.

Notes

- (m) The Receiving Party’s outside counsel of record and/or authorized consultants and experts pursuant to paragraph 5(g) above shall be entitled to take handwritten notes relating to the Source Code Material including identifying the file name and line number of the code, but may not copy entire lines of code. Such notes shall be labeled “RESTRICTED CONFIDENTIAL - SOURCE CODE.” The Receiving Party shall be under no obligation to produce handwritten notes. No copies of all or any portion of the Source Code Material may leave the source code review room. The Receiving Party may take handwritten notes. All notes are protected by attorney client privilege. No other written or electronic record of the Source Code Material is permitted except as otherwise provided in this Protective Order.
- (n) Unless otherwise agreed in advance by the parties in writing, the Receiving Party’s authorized counsel and/or authorized consultants and experts pursuant to paragraph 5(g) above shall remove all notes, documents, and all other materials from the source code review room that may contain work product and/or attorney-client privileged information at the end of each day. Materials inadvertently left in the source code review room do not operate as a waiver of the attorney work product doctrine or any other applicable privilege and shall be returned to the owner promptly. The Producing Party shall not be responsible for any items left in the source code review room.

Printouts and Copies

- (o) The Receiving Party may request a printout of reasonable portions of the Source Code Material. The Receiving Party shall not request the printouts of Source Code Material in order to review blocks of Source Code Material in the first instance (*i.e.*, as an alternative to reviewing Source Code Material electronically on the source code computers) as the Parties acknowledge and agree that the purpose of the protections herein would be frustrated by printing portions of Source Code Material for review and analysis elsewhere.
- (p) In the event that the Producing Party objects that identified files are not reasonably necessary, the Producing Party shall make such objection known to the Receiving

Party within three (3) business days of the identification of any files by the Receiving Party. The Producing Party will make itself available within three (3) business days of making the objection for a meet and confer. If, after meeting and conferring, the Producing Party and the Receiving Party cannot resolve the objection, the Producing Party shall seek a judicial resolution of whether or not the identified Source Code Material in question is reasonably necessary to any case preparation activity. The parties recognize that the volume and format of source code at issue in the case is unknown. Nevertheless, the Receiving Party will make a good faith effort to not exceed the agreed limit of fifteen (15) consecutive pages and an aggregate of more than six-hundred fifty (650) pages of printouts of Source Code Material. Requests to print additional pages beyond these limits will not be unreasonably withheld by the Producing Party. In the event Producing Party refuses a request to print additional pages, the Receiving Party will bear the burden of demonstrating the additional printouts are relevant to the needs of the case. At the request of the Receiving Party, in the absence of any objection the Producing Party shall provide within five (5) days of such request, a printout of the Source Code Material requested by the Receiving Party. The Producing Party shall print and label such Source Code Material with the complete path name, if applicable, and append production numbers and the designation "RESTRICTED CONFIDENTIAL - SOURCE CODE."

- (q) Unless agreed by the parties, images or copies of Source Code Material shall not be included in correspondence between the Parties (references to production numbers shall be used instead). The parties will make a good faith effort to omit or minimize the use of images or copies of Source Code Material in pleadings, exhibits, and other papers.
- (r) The Receiving Party shall maintain a record of any individual who has inspected any portion of the Source Code Material in electronic or paper form. The log shall include the number of paper copies, names of the recipients of paper copies, and locations where the paper copies are stored. Upon three (3) business day's advance notice to the Receiving Party by the Producing Party, the Receiving Party shall provide a copy of this log to the Producing Party. The Parties agree the purpose of the log is to evidence compliance with this Order and that the log may not be used for examination of a witness. All copies of Source Code Material shall be printed with an "anti-copying" watermark that reveals whether such printouts have been photocopied.
- (s) Consistent with and subject to paragraph 10(k) above, Source Code Material must be kept at all times in a secured, locked and hardened container (e.g., a safe or lockbox) that prevents duplication of or unauthorized access to the Source Code Material.

Depositions

- (t) For all depositions, Receiving Party may request that the witness have access to a

Secure Source Code Computer during the deposition three (3) days prior to the deposition. Upon such request, Producing Party will ensure a Secure Source Code Computer is available with all code during examination with reasonable search tools so that the witness can be questioned about the code. If a Source Code Material document has been marked up or altered in any way by the deponent, the Producing Party shall copy the exhibit in its marked form and provide one copy to outside counsel for the Receiving Party. Copies of Source Code Material that are marked as deposition exhibits shall not be provided to the Court Reporter or attached to deposition transcripts; rather, the deposition record will identify the exhibit by its production numbers.

11. Any attorney representing a Party, whether in-house or outside counsel, and any person associated with a Party that receives the other Party's Protected Material that is designated RESTRICTED - ATTORNEYS' EYES ONLY and/or RESTRICTED CONFIDENTIAL - SOURCE CODE (collectively "HIGHLY SENSITIVE MATERIAL"), who obtains, receives, has access to, or otherwise learns, in whole or in part, HIGHLY SENSITIVE MATERIAL under this Order shall not prepare, prosecute, supervise, or assist in the preparation or prosecution of any patent application pertaining to the field of the invention of the patents-in-suit on behalf of the Receiving Party or its acquirer, successor, predecessor, or other affiliate during the pendency of this Action and for one year after its conclusion, including any appeals. This provision does not prohibit a Party's counsel of record or experts in this litigation from participating in or representing it in reexamination proceedings, Post-Grant Review proceedings, Inter Partes Review proceedings, or Covered Business Method Review proceedings involving any patent, including the patents-in-suit, provided they do not advise on, consult on, prepare, draft, edit, or otherwise provide input into any claims, or amendment to specifications or claims in those proceedings. To ensure compliance with the purpose of this provision, each Party shall create an "Ethical Wall" between those persons with access to HIGHLY SENSITIVE MATERIAL and any individuals who, on behalf of the Party or its acquirer, successor,

predecessor, or other affiliate, prepare, prosecute, supervise or assist in the preparation or prosecution of any patent application pertaining to the field of invention of the patent-in-suit.

12. Unless otherwise provided by this Order or provided by the Court, confidentiality obligations imposed by this Order shall remain in effect following the conclusion of this lawsuit.
13. Nothing in this Order shall require production of documents, information or other material that a Party contends is protected from disclosure by the attorney-client privilege, the work product doctrine, or other privilege, doctrine, or immunity. If documents, information or other material subject to a claim of attorney-client privilege, work product doctrine, or other privilege, doctrine, or immunity is inadvertently or unintentionally produced, such production shall in no way prejudice or otherwise constitute a waiver of, or estoppel as to, any such privilege, doctrine, or immunity. Any Party that inadvertently or unintentionally produces documents, information or other material it reasonably believes are protected under the attorney-client privilege, work product doctrine, or other privilege, doctrine, or immunity may obtain the return of such documents, information or other material by promptly notifying the recipient(s) and providing a privilege log for the inadvertently or unintentionally produced documents, information or other material. The recipient(s) shall gather and return all copies of such documents, information or other material to the Producing Party, except for any pages containing privileged or otherwise protected markings by the recipient(s), which pages shall instead be destroyed and certified as such to the Producing Party. If a Receiving Party believes that a produced document or other information may reasonably be subject to a claim of privilege, then the Receiving Party

shall promptly inform the Producing Party of the beginning BATES number of the document or, if no BATES number is available, shall otherwise inform the Producing Party of the information.

14. There shall be no disclosure of any DESIGNATED MATERIAL by any person authorized to have access thereto to any person who is not authorized for such access under this Order. The Parties are hereby ORDERED to safeguard all such documents, information and material to protect against disclosure to any unauthorized persons or entities.
15. Nothing contained herein shall be construed to prejudice any Party's right to use any DESIGNATED MATERIAL in taking testimony at any deposition or hearing provided that the DESIGNATED MATERIAL is only disclosed to a person(s) who is: (i) eligible to have access to the DESIGNATED MATERIAL by virtue of his or her employment with the designating party, (ii) identified in the DESIGNATED MATERIAL as an author, addressee, or copy recipient of such information, (iii) although not identified as an author, addressee, or copy recipient of such DESIGNATED MATERIAL, has, in the ordinary course of business, seen such DESIGNATED MATERIAL, (iv) a current or former officer, director or employee of the Producing Party or a current or former officer, director or employee of a company affiliated with the Producing Party; (v) counsel for a Party, including outside counsel and in-house counsel (subject to paragraph 9 of this Order); (vi) an independent contractor, consultant, and/or expert retained for the purpose of this litigation; (vii) court reporters and videographers; (viii) the Court; or (ix) other persons entitled hereunder to access to DESIGNATED MATERIAL. DESIGNATED MATERIAL shall not be disclosed to any other persons unless prior authorization is obtained from counsel representing the Producing Party or from the Court.

16. Parties may, at the deposition or hearing or within thirty (30) days after receipt of a deposition or hearing transcript, designate the deposition or hearing transcript or any portion thereof as “CONFIDENTIAL,” “RESTRICTED - ATTORNEY’ EYES ONLY,” or “RESTRICTED CONFIDENTIAL - SOURCE CODE” pursuant to this Order. Access to the deposition or hearing transcript so designated shall be limited in accordance with the terms of this Order. Until expiration of the 30-day period, the entire deposition or hearing transcript shall be treated as confidential.
17. Any DESIGNATED MATERIAL that is filed with the Court shall be filed under seal and shall remain under seal until further order of the Court. The filing party shall be responsible for informing the Clerk of the Court that the filing should be sealed and for placing the legend “FILED UNDER SEAL PURSUANT TO PROTECTIVE ORDER” above the caption and conspicuously on each page of the filing. Exhibits to a filing shall conform to the labeling requirements set forth in this Order. If a pretrial pleading filed with the Court, or an exhibit thereto, discloses or relies on confidential documents, information or material, such confidential portions shall be redacted to the extent necessary and the pleading or exhibit filed publicly with the Court.
18. The Order applies to pretrial discovery. Nothing in this Order shall be deemed to prevent the Parties from introducing any DESIGNATED MATERIAL into evidence at the trial of this Action, or from using any information contained in DESIGNATED MATERIAL at the trial of this Action, subject to any pretrial order issued by this Court.
19. A Party may request in writing to the other Party that the designation given to any DESIGNATED MATERIAL be modified or withdrawn. If the designating Party does not agree to redesignation within ten (10) days of receipt of the written request, the requesting

Party may apply to the Court for relief. Upon any such application to the Court, the burden shall be on the designating Party to show why its classification is proper. Such application shall be treated procedurally as a motion to compel pursuant to Federal Rules of Civil Procedure 37, subject to the Rule's provisions relating to sanctions. In making such application, the requirements of the Federal Rules of Civil Procedure and the Local Rules of the Court shall be met. Pending the Court's determination of the application, the designation of the designating Party shall be maintained.

20. Each outside consultant or expert to whom DESIGNATED MATERIAL is disclosed in accordance with the terms of this Order shall be advised by counsel of the terms of this Order, shall be informed that he or she is subject to the terms and conditions of this Order, and shall sign an acknowledgment that he or she has received a copy of, has read, and has agreed to be bound by this Order. A copy of the acknowledgment form is attached as Appendix A.
21. To the extent that any discovery is taken of persons who are not Parties to this Action ("Third Parties") and in the event that such Third Parties contended the discovery sought involves trade secrets, confidential business information, or other proprietary information, then such Third Parties may agree to be bound by this Order.
22. To the extent that discovery or testimony is taken of Third Parties, the Third Parties may designate as "CONFIDENTIAL" or "RESTRICTED - ATTORNEYS' EYES ONLY" any documents, information or other material, in whole or in part, given by such Third Parties. The Third Parties shall have ten (10) days after production of such documents, information or other materials to make such a designation. Until that time period lapses or until such a designation has been made, whichever occurs sooner, all documents, information or other

material so produced or given shall be treated as “CONFIDENTIAL” in accordance with this Order.

23. Within sixty (60) days of final termination of this Action, including any appeals, all DESIGNATED MATERIAL, including all copies, duplicates, abstracts, indexes, summaries, descriptions, and excerpts or extracts thereof (excluding excerpts or extracts incorporated into any privileged memoranda of the Parties and materials which have been admitted into evidence in this Action), shall at the Producing Party’s election either be returned to the Producing Party or be destroyed. No copies shall be retained by a Receiving Party or other individuals retained to assist them in this litigation, except counsel may retain copies of pleadings in their files and the terms of this Protective Order shall continue to apply to such copies. The language contained in this paragraph shall not preclude counsel from retaining documents constituting attorney work product as defined by applicable law, or from retaining one copy of all such work product documents as part of a permanent litigation file that is otherwise subject to the confidentiality restrictions set forth herein. The Receiving Party shall verify the return or destruction by affidavit furnished to the Producing Party, upon the Producing Party’s request. Nothing in this order requires a Party to destroy any information it is required by law to retain. The Parties shall not be required to delete information that may reside on electronic disaster recovery systems that are overwritten in the normal course of business, or information that may reside in electronic files which are not reasonably accessible. However, parties and their Counsel shall not retrieve, access, nor use any Protected Material from said electronic disaster recovery systems or other not-reasonably accessible data sources after the conclusion of the matter. The parties shall make a good faith effort to locate and destroy any case materials

designated RESTRICTED CONFIDENTIAL - SOURCE CODE.

24. The failure to designate documents, information or material in accordance with this Order and the failure to object to a designation at a given time shall not preclude the filing of a motion at a later date seeking to impose such designation or challenging the propriety thereof. The entry of this Order and/or the production of documents, information and material hereunder shall in no way constitute a waiver of any objection to the furnishing thereof, all such objections being hereby preserved.
25. Any Party knowing or believing that any other party is in violation of or intends to violate this Order and has raised the question of violation or potential violation with the opposing party and has been unable to resolve the matter by agreement may move the Court for such relief as may be appropriate in the circumstances. Pending disposition of the motion by the Court, the Party alleged to be in violation of or intending to violate this Order shall discontinue the performance of and/or shall not undertake the further performance of any action alleged to constitute a violation of this Order.
26. Production of DESIGNATED MATERIAL by each of the Parties shall not be deemed a publication of the documents, information and material (or the contents thereof) produced so as to void or make voidable whatever claim the Parties may have as to the proprietary and confidential nature of the documents, information or other material or its contents.
27. The Receiving Party may not upload to, share with, or otherwise use DESIGNATED MATERIAL of the Producing Party on any generative artificial intelligence (GenAI) or large language model (LLM) tools or platforms including any third-party, public, or internal tools unless the Producing Party approves the use of such tool or the Court issues an Order approving the use of such tool upon request by the Receiving Party.

28. Nothing in this Order shall be construed to effect an abrogation, waiver or limitation of any kind on the rights of each of the Parties to assert any applicable discovery or trial privilege.
29. Each of the Parties shall also retain the right to file a motion with the Court (a) to modify this Order to allow disclosure of DESIGNATED MATERIAL to additional persons or entities if reasonably necessary to prepare and present this Action and (b) to apply for additional protection of DESIGNATED MATERIAL.

So ORDERED and SIGNED this 21st day of November, 2025.



RODNEY GILSTRAP
UNITED STATES DISTRICT JUDGE

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

PERCEPTIVE AUTOMATA LLC,
Plaintiff,

v.

TESLA, INC.,
Defendant.

§
§
§
§
§
§
§
§

Case No. 2:25-cv-00742-JRG

JURY TRIAL DEMANDED

**APPENDIX A
UNDERTAKING OF EXPERTS OR CONSULTANTS REGARDING
PROTECTIVE ORDER**

I, _____, declare that:

1. My address is _____.
My current employer is _____.
My current occupation is _____.
2. I have received a copy of the Protective Order in this action. I have carefully read and understand the provisions of the Protective Order.
3. I will comply with all of the provisions of the Protective Order. I will hold in confidence, will not disclose to anyone not qualified under the Protective Order, and will use only for purposes of this action any information designated as “CONFIDENTIAL,” “RESTRICTED - ATTORNEYS’ EYES ONLY,” or “RESTRICTED CONFIDENTIAL - SOURCE CODE” that is disclosed to me.
4. Promptly upon termination of these actions, I will return all documents and things designated as “CONFIDENTIAL,” “RESTRICTED - ATTORNEYS’ EYES ONLY,” or “RESTRICTED CONFIDENTIAL - SOURCE CODE” that came into my possession,

and all documents and things that I have prepared relating thereto, to the outside counsel for the party by whom I am employed.

5. I hereby submit to the jurisdiction of this Court for the purpose of enforcement of the Protective Order in this action.

I declare under penalty of perjury that the foregoing is true and correct.

Signature _____

Date _____

Attachment C

Motion Prediction of Traffic Actors for Autonomous Driving using Deep Convolutional Networks

Nemanja Djuric, Vladan Radosavljevic, Henggang Cui,
Thi Nguyen, Fang-Chieh Chou, Tsung-Han Lin, Jeff Schneider
Uber Advanced Technologies Group
{ndjuric, vradosavljevic, hcui2, thi, fchou, hanklin, jschneider}@uber.com

Abstract

Recent algorithmic improvements and hardware breakthroughs resulted in a number of success stories in the field of AI impacting our daily lives. However, despite its ubiquity AI is only just starting to make advances in what may arguably have the largest impact thus far, the nascent field of autonomous driving. In this work we discuss this important topic and address one of crucial aspects of the emerging area, the problem of predicting future state of autonomous vehicle's surrounding necessary for safe and efficient operations. We introduce a deep learning-based approach that takes into account current state of traffic actors and produces rasterized representations of each actor's vicinity. The raster images are then used by deep convolutional models to infer future movement of actors while accounting for inherent uncertainty of the prediction task. Extensive experiments on real-world data strongly suggest benefits of the proposed approach. Moreover, following successful tests the system was deployed to a fleet of autonomous vehicles.

1 Introduction

Recent years have witnessed unprecedented advances in Artificial Intelligence (AI) applications, with a number of highly publicized success stories that have captured attention and imagination of professionals and laymen alike. Computers have reached and even surpassed human performance in centuries old games such as go and chess [33, 34], are starting to understand health conditions and suggest medical treatment [38], and can even reason about complex relationships conveyed through images [39]. Prior to this, AI was already well-established at the very core of multibillion-dollar industries such as advertising [5] and finances [26], where algorithmic approaches have all but replaced human experts. Nevertheless, despite huge strides the AI revolution is far from over [12]. One of the major industries mostly untouched by the ongoing progress is the automobile domain, where AI may potentially make its greatest impact yet through the advent of self-driving cars.

Driving a motor vehicle is a complex undertaking, requiring drivers to understand involved multi-actor scenes in real time and act upon rapidly changing environment within a fraction of a second (*actor* is a generic term referring to any surrounding vehicle, pedestrian, bicycle, or other potentially moving object). Unfortunately, humans are famously ill-fitted for the task. This is sadly corroborated by grim road statistics that often worsen year after year, showing that traffic accidents were the number four cause of death in the US in 2015, accounting for more than 5% of the total [24]. In addition, despite large investments by governments and progress made in traffic safety technologies, 2017 was still one of the deadliest years for motorists in the past decade [25]. Moreover, human error is considered responsible for up to 94% of car crashes [36], strongly suggesting that mitigating or even removing the unreliable human factor could potentially save hundreds of thousands of motorists' lives and tens of billions of dollars in accident-related damages and medical expenses [4].




Figure 1: Example of a complex scene handled by our model; (a) scene in a 3D viewer, with indicated actor of interest; (b) rasterized surrounding in bird’s eye view of the actor of interest (colored red) used as an input to CNN, with overlaid trajectories (dotted green line - ground truth; blue - predicted)

Latest advances in AI and breakthroughs in high-performance computing, allowing for more powerful processing hardware at lower costs, unlocked the potential to reverse the regrettable negative trend of safety on our public roads. In particular, together they gave rise to development of self-driving technology, where driving decisions are entrusted to a computer aboard an autonomous vehicle (AV), equipped with a number of external sensors and capable of processing large amounts of information at speeds and throughputs far surpassing human capabilities. Once fully matured the technology is expected to drastically improve road safety and redefine the very way we organize transportation and even our lives [29]. To this end, the industry and governments are working together to fulfill this potential and bring the AVs to consumers, with front-runners such as Waymo, Uber, and Ford investing significant resources into autonomous research, and states such as Arizona, Nevada, and California enacting necessary legal frameworks. Nevertheless, autonomous driving is still in its initial phases of development, with a number of challenges lying ahead of the researchers.

To safely deploy AVs to public roads one must solve a sequence of tasks that include detection and tracking of actors in AV’s surrounding, predicting their future trajectories, as well as navigating the AV safely and effectively towards its intended destination while taking into account current and future states of the actors. In this work we focus on a critical component of this pipeline, predicting future trajectory of tracked vehicles (in the following we use the terms *vehicle* and *actor* interchangeably), where a working detection and tracking system is assumed. Our main contributions are as follows:

- We propose to rasterize surrounding of each vehicle in AV’s vicinity, thus providing full context and information necessary for accurate future trajectory prediction;
- We trained deep convolutional neural network (CNN) to predict future vehicle trajectories, while accounting for inherent uncertainty of vehicle motion in road traffic;
- We run evaluation on large-scale real-world data showing that the system provides accurate predictions and well-calibrated uncertainties, clearly indicating its practical benefits;
- Following extensive testing in both offline and online setting, the system was successfully deployed to a fleet of self-driving vehicles.

Example of a complex scene handled by our model is shown in Figure 1. Figure 1(a) shows the scene in an internal 3D viewer, while Figure 1(b) shows the corresponding rasterized 2D input image (or *raster*) and the predicted trajectory (we refer to the actor being considered and whose surrounding corresponds to the input raster as *actor of interest*). The method uses rasterized vehicle context (including map and other actors) to accurately predict actor’s movement in a dynamic environment.

2 Related work

In the past decade a number of methods were proposed to tackle the problem of predicting future motion of traffic actors. Comprehensive overview of the topic can be found in [23, 43]. Here, we review literature from the perspective of autonomous driving. We first cover engineered approaches commonly used in practice in the self-driving industry. Then, we discuss learned approaches using classical machine learning methods as well as deep learning.

2.1 Motion prediction in self-driving systems

Accurate prediction of actors' motion is a critical component of an effective self-driving system [7, 46]. In particular, prediction is tightly coupled with AV's egomotion planning, as it is essential to accurately estimate future world state in order to correctly plan AV's path through a highly dynamic environment. Inaccurate motion prediction may lead to accidents, as exemplified by a collision between MIT's "Talos" and Cornell's "Skyne" during the 2007 DARPA Urban Challenge [8].

Most of the deployed self-driving systems use well-established engineered approaches for motion prediction. The common approach consists of computing object's future motion by propagating its state over time based on kinematics and assumptions of underlying physical system. State estimate usually comprises position, speed, acceleration, and heading of the object, and techniques such as Kalman filter (KF) [15] are used to estimate and propagate the state in the future. For example, in Honda's deployed system [7], KF tracker is used to predict motion of vehicles around AV. While this approach works well for short-term predictions, its performance degrades for longer horizons as the model ignores surrounding context (e.g., roads, other actors, traffic rules). On the other hand, Mercedes Benz's motion prediction component uses map information as a constraint to compute vehicle's future position [46]. The system first associates each detected vehicle with one or more lanes from the map. Then, all possible paths are generated for each (*vehicle, associated lane*) pair based on map topology, lane connectivity, and vehicle's current state estimate. This heuristic provides reasonable predictions in common cases, however it does not scale well and is not able to model unusual traffic scenarios. As an alternative to existing deployed engineered approaches, by considering large amounts of training data our proposed approach automatically learns the fact that vehicles usually obey road and lane constraints, while being capable of also handling outliers.

2.2 Learned prediction models

Manually designed engineered models often impose unrealistic assumptions not supported by the data (e.g., that traffic always follows designated lanes), which motivated use of learned models as an alternative. A large class of learned models are maneuver-based models (e.g., using Hidden Markov Model [37]) which are object-centric approaches that predict discrete action of each object independently. The independence assumption does not often hold true, which is mitigated by the use of Bayesian networks [31] that are computationally more expensive and not feasible in real-time tasks. Additionally, in [3] authors learned scene-specific motion patterns and applied them to novel scenes with an image-based similarity function. However, these methods also require manually designed features to capture context information, resulting in suboptimal performance. Alternatively, Gaussian Process (GP) regression can be used to address the motion prediction problem [41]. GP regression is well-suited for the task with desirable properties such as ability to quantify uncertainty, yet it is limited when modeling complex actor-environment interactions. In recent work researchers focused on how to model environmental context using Inverse Reinforcement Learning (IRL) [45]. Kitani *et al.* [18] used inverse optimal control to predict pedestrian paths by considering scene semantics, however the proposed IRL methods are inefficient for real-time applications.

The success of deep learning [11] motivated its use for motion prediction. In [22] a Recurrent Neural Network (RNN)-based method was proposed for long-term predictions of multiple interacting agents given scene context. In [2] authors proposed a social Long Short-Term Memory (LSTM) to model human movement together with social interactions. Authors of [9] used LSTM to predict ball motion in billiards using a sequence of visual glimpses. In [42] LSTMs were used to classify basketball plays, with overhead raster images as inputs. Similarly, the authors of [27, 28] used overhead raster as input to RNNs to predict raster image of a scene in a next timestep, unlike our work where per-object trajectories are directly inferred. Due to strict time constraints of a deployed real-time system, as well as a requirement to more easily debug and understand model decisions made on public roads, we used simpler feed-forward CNN architectures for the task of motion prediction.

3 Proposed approach

Let us assume that we have access to real-time data streams coming from sensors such as lidar, radar, or camera, installed aboard a self-driving vehicle. Furthermore, we assume to have an already functioning tracking system ingesting the sensor data, allowing detection and tracking of traffic actors in real-time. For example, we can make use of any of a number of Kalman filter-based methods that

have found wide practical use [6], taking sensor data as input and outputting tracks of individual actors that represent their state estimates at fixed intervals. State estimates contain the following information describing an actor: bounding box, position, velocity, acceleration, heading, and heading change rate. Lastly, we assume access to mapping data of an area where the AV is operating, comprising road and crosswalk locations, lane directions, and other relevant map information.

Let us denote overall map data by $\mathcal{M}$, and a set of discrete times at which tracker outputs state estimates as $\mathcal{T} = \{t_1, t_2, \dots, t_T\}$, where time gap between consecutive time steps is constant (e.g., gap is equal to $0.1s$ for tracker running at frequency of $10Hz$). Then, we denote state output of a tracker for the i -th actor at time t_j as $\mathbf{s}_{ij}$, where $i = 1, \dots, N_j$ with N_j being a number of unique actors tracked at t_j . Note that in general the actor counts vary for different time steps as new actors appear within and existing ones disappear from the sensor range. Then, given data $\mathcal{M}$ and all actors' state estimates up to and including time step t_j (denoted by $\mathcal{S}_j$), the task is to predict sequence of future states $[\mathbf{s}_{i(j+1)}, \dots, \mathbf{s}_{i(j+H)}]$, where H denotes the number of future consecutive time steps for which we predict states (or *prediction horizon*). Without the loss of generality, in this work we simplify the task to infer i -th actor's future positions instead of full state estimates, denoted as $[x_{i(j+1)}, \dots, x_{i(j+H)}]$ for x -positions and $[y_{i(j+1)}, \dots, y_{i(j+H)}]$ for y -positions. Past and future positions at time t_j are represented in actor-centric coordinate system derived from actor's state at time t_j , where forward direction represents x -axis, left-hand direction represents y -axis, and actor's bounding box centroid represents the origin.

3.1 Model inputs

To model dynamic actor context at time step t_j we use state data $\mathcal{S}_j$, comprising actors' states until t_j , inclusive. In addition, to model static context we use map data $\mathcal{M}$, comprising road and crosswalk polygons, as well as lane directions and boundaries. Road polygons describe drivable surface, lane describes driving path for a single vehicle, while crosswalk polygons describe road surface used for pedestrian crossing. Lanes are represented by boundaries and directed lines positioned at the center.

Instead of manually defining features that represent context for each individual actor, we propose to rasterize scene for the i -th actor at time step t_j into an RGB image (see Figure 1 for example). Then, we train CNN using rasterized images as inputs to predict actor's trajectory, where the network automatically infers relevant features. Optionally, the model can also take as input a current state of the actor of interest $\mathbf{s}_{ij}$ represented as a vector (see Section 3.3 for details of the architecture).

3.1.1 Rasterization

To describe rasterization, let us first introduce a concept of *vector layer*, formed by a collection of polygons and lines that belong to a common type. For example, in the case of map elements we have vector layer of roads, of crosswalks, and so on. To rasterize vector layer into an RGB space, each vector layer is manually assigned a color from a set of distinct RGB colors that make a difference among layers more prominent. The only layer that does not have its defined RGB color is a layer that encodes lane direction. Instead of assigning a specific RGB color, we use a direction of each straight line segment as hue value in HSV color space, with saturation and value set to maximum. We then convert HSV to RGB color space, thus encoding driving direction of each lane in the resulting rasterized image. For example, in Figure 1 we can see that lanes going in opposite directions are represented by colors diametrically opposite to each other on the HSV color cylinder. Once the colors are defined, vector layers are rasterized one by one on top of each other, in the order from layers that represent larger areas such as road polygons towards layers that represent finer structures such as lanes or actors' bounding boxes. Important rasterization parameter is pixel resolution, which we set to $0.1m$ considering the trade-off between image size and ability to represent fine details.

As discussed earlier, we are interested in representing context around each actor separately. To represent context around the i -th actor tracked at time step t_j we create a rasterized image I_{ij} of size $n \times n$ such that the actor is positioned at pixel (w, h) within I_{ij} , where w represents width and h height measured from the bottom-left corner of the image. The image is rotated such that actor's heading points up, where lane directions are computed relative to the actor's heading and then encoded in the HSV space. We set $n = 300$, actor of interest is positioned at $w = 150$ and $h = 50$, so that $25m$ of context in front of the actor and $5m$ from the back is taken into account. Lastly, we color the actor of interest differently so that the vehicle for which we predict trajectory is distinguishable from other vehicles (as seen in Figure 1 it is colored red while other actors are colored yellow).




Figure 2: The proposed network architecture combining raster and state inputs

To represent actors' histories, their bounding boxes at consecutive time steps $[t_{j-K+1}, \dots, t_j]$ are rasterized in the raster image on top of map vector layers. Each historical actor polygon is rasterized with the same color as the current polygon yet with reduced level of brightness, resulting in the *fading* effect. Brightness level at t_{j-k} is equal to $\max(0, 1 - k \cdot \delta)$, $k = 0, 1, \dots, K - 1$, where we set $\delta = 0.1$ and K to either 1 (no fading) or 5 (with fading, example shown in Figure 1).

3.2 Optimization problem

To obtain analytical expressions for loss functions used to optimize deep networks, let us first introduce *displacement error* for the i -th actor at time t_j for horizon $h \in \{1, \dots, H\}$,

$$d_{i(j+h)} = \sqrt{(x_{i(j+h)} - \hat{x}_{i(j+h)}(\mathcal{S}_j, \mathcal{M}, \theta))^2 + (y_{i(j+h)} - \hat{y}_{i(j+h)}(\mathcal{S}_j, \mathcal{M}, \theta))^2}, \quad (1)$$

defined as Euclidean distance between observed and predicted positions. Here, θ denotes parameters of a model, while $\hat{x}_{i(j+h)}(\mathcal{S}_j, \mathcal{M}, \theta)$ and $\hat{y}_{i(j+h)}(\mathcal{S}_j, \mathcal{M}, \theta)$ denote position outputs of the model that takes available states $\mathcal{S}_j$ and map $\mathcal{M}$ as inputs. Then, overall loss incurred by predicting trajectory for a complete prediction horizon is equal to average squared displacement error of trajectory points,

$$L_{ij} = \frac{1}{H} \sum_{h=1}^H d_{i(j+h)}^2, \quad (2)$$

where we train the model to output $2H$ -dimensional vector, representing predicted x - and y -positions for each of H trajectory points. Optimizing over all actors and time steps found in the training data, we find optimal parameters by minimizing overall training loss,

$$\theta^* = \arg \min_{\theta} \mathcal{L} = \arg \min_{\theta} \sum_{j=1}^T \sum_{i=1}^{N_j} L_{ij}. \quad (3)$$

Alternatively, as the prediction task is inherently noisy it is useful to capture aleatoric uncertainty present in the data [16, 21], in addition to optimizing for point estimate as in (3). To that end, we assume that displacement errors are sampled from a half-normal distribution [14], denoted as

$$d_{i(j+h)} \sim \mathcal{FN}(0, \hat{\sigma}_{i(j+h)}(\mathcal{S}_j, \mathcal{M}, \theta)^2), \quad (4)$$

where standard deviation $\hat{\sigma}_{i(j+h)}$ is computed by the model. Then, we can write overall loss for the i -th actor at time t_j as negative log-likelihood of the observed data, equal to

$$L_{ij} = \sum_{h=1}^H \left(\frac{d_{i(j+h)}^2}{2 \hat{\sigma}_{i(j+h)}(\mathcal{S}_j, \mathcal{M}, \theta)^2} + \log \hat{\sigma}_{i(j+h)}(\mathcal{S}_j, \mathcal{M}, \theta) \right), \quad (5)$$

where we train the model to output $3H$ -dimensional vector, representing predicted x -position, y -position, and standard deviation of error for each of H trajectory points. Lastly, optimizing over entire training data we solve (3) with L_{ij} computed as in (5).

3.3 Network architecture

In this section we describe an architecture used to solve the optimization problems (2) and (5), which we also illustrate in Figure 2. To extract features from an input raster we can use any existing CNN architecture (referred to as *base CNN*). In addition, to input actor state we encode it as a 3-D vector comprising velocity, acceleration, and heading change rate (position and heading are not required as they were already used during raster generation), and concatenate the resulting vector with flattened output of the base CNN. Then, the combined features are passed through a fully-connected layer (we set its size to 4,096) connected to an output layer of size $2H$ if solving (2), or $3H$ if solving (5).

Table 1: Comparison of average prediction errors for competing methods (in meters)

Predictor	Raster	State	Loss	Displacement	Along-track	Cross-track
UKF	–	–	–	2.36	2.69	0.49
AlexNet	w/o fading	no	(2)	3.14	3.11	0.35
AlexNet	w/ fading	no	(2)	1.24	1.23	0.22
AlexNet	w/o fading	yes	(2)	0.97	0.94	0.21
AlexNet	w/ fading	yes	(2)	0.86	0.83	0.20
VGG-19	w/ fading	yes	(2)	0.77	0.75	0.19
ResNet-50	w/ fading	yes	(2)	0.76	0.74	0.18
MobileNet-v2	w/ fading	yes	(2)	0.73	0.70	0.18
MobileNet-v2	w/ fading	yes	(5)	0.71	0.68	0.18

4 Experiments

Baseline We used Unscented Kalman filter (UKF) [40] with dynamic vehicle model [19] as a strong non-CNN baseline, taking raw sensor data from camera, lidar, and radar and outputting state estimates for each tracked vehicle. The tracking filter was readily available as it is a default tracker on our fleet, is pretrained on large amount of labeled data, is highly optimized and successfully tested on millions on driven miles (unfortunately, no other details can be given due to confidentiality concerns). Then, we can use this tracker to predict future motion by forward propagating estimated states in time.

Data We collected 240 hours of data by manually driving AV in Pittsburgh, PA and Phoenix, AZ in various traffic conditions (e.g., varying times of day, days of the week), with data collection rate of $10Hz$ (the same frequency the UKF was run on). Each tracked actor at each discrete tracking time step amounts to a single data point, with overall data comprising 7.8 million data points. We considered prediction horizon of $3s$ (i.e., we set $H = 30$), and used 3:1:1 split to obtain train/validation/test data.

Models We compared the baseline to several variants of the proposed approach. We trained and evaluated the following base CNNs: AlexNet [20], VGG-19 [35], ResNet-50 [13], and MobileNet-v2 (MNv2) [30]. Furthermore, to evaluate how varying input complexity affects performance we considered architectures that use: 1) raster without fading and state, solving (2); 2) raster with fading and without state, solving (2); 3) raster without fading and with state, solving (2); 4) raster with fading and state, solving (2); 5) raster with fading and state, and outputting uncertainty, solving (5).

Training Models were implemented in TensorFlow [1] and trained on 16 Nvidia Titan Xs connected using commodity network hardware. To coordinate the GPUs we used open-source distributed framework Horovod [32], completing training in around 24 hours. We used batch size of 32 and trained with Adam optimizer [17], where we set initial learning rate to 10^{-4} that is further decreased by a factor of 0.9 every 20 thousand iterations. All models were trained end-to-end from scratch, except for a model with uncertainty outputs which was initialized with a corresponding model without uncertainty and then fine-tuned (training from scratch did not give satisfactory results). We deployed models to an AV with an undisclosed GPU aboard, performing batch inference in $10ms$ on average.

4.1 Results

In Table 1 we report error metrics averaged over the prediction horizon that are relevant for motion prediction: displacement errors, as well as along-track and cross-track errors [10].

We first conducted an ablation study using AlexNet, performing a sequence of experiments with varying input complexity (given in the upper half of Table 1). When we provide neither fading nor state inputs the model performs worse than UKF, as the network does not have enough information to estimate current state of an actor from the raster itself. Interestingly, when we include fading the model starts to outperform the baseline by a large margin, indicating that actor’s state can be inferred solely from providing past positions through fading. If instead of fading we directly provide state estimates we get even better performance, as the state info is already distilled and does not have to be estimated from raster. Furthermore, using raster with fading together with state inputs leads to additional performance boost, suggesting that fading carries additional information not available through state itself, and that raster image and other external inputs can be seamlessly combined through the proposed architecture to obtain improved accuracy.




Figure 3: Reliability diagrams at the prediction horizon of: (a) 1 second; (b) 3 seconds

Next, we run experiments with several popular CNN architectures as base CNNs. We found that VGG and ResNet provide improvements over the baseline AlexNet, as observed in other applications as well [35]. However, both are outperformed by the novel MNv2 architecture that combines a number of deep learning ideas under one roof (e.g., bottleneck layers, residual connections, depthwise convolutions). Lastly, taking the best performing MNv2 as base and adding uncertainty outputs led to the best results. Not only do additional outputs allow estimation of trajectory uncertainty in addition to trajectory point estimates, but they also help mitigate the adverse effect of noisy data points during training. In the following we analyze results of this best performing model in greater detail.

We used reliability diagrams to evaluate how closely predicted error distribution matches testing error distribution. The diagrams are generated by measuring how large is an observed displacement error compared to a predicted confidence, and computing what fraction of observed errors falls within the expected range given by the estimated standard deviation. For example, due to the assumption of Gaussianity we expect 68% of observed errors to be within the predicted one sigma, and diagram point at predicted value of 0.68 should be as close as possible to observed value of 0.68. Thus, closer the curve is to the diagonal line, better calibrated is the model. Figure 3 shows the diagrams for prediction horizons of 1s and 3s. The predicted uncertainties are well aligned with the reference line, especially for the 3s-predictions while 1s-predictions are slightly underconfident. Thus, given an estimated sigma output, we can expect with high confidence that in 68% of cases the actual error will not be larger than that value. Plots for other horizons are omitted as they resemble the ones shown.

4.2 Case studies

In the three rows of Figure 4 we give example outputs for three interesting scenes commonly encountered in traffic. As we will see, in all three cases the model provided accurate short-term trajectories as well as reasonable and intuitive uncertainty estimates.

The first case (shown in the first row) involves actor cutting over opposite lane when entering from off-street parking, where the model correctly predicts that the actor will queue for vehicles in front (see the image in the first column). The uncertainty estimates reflect peculiarity of the situation (see the image in the second column), as the actor is not following common traffic rules and may choose to either queue for the leftmost vehicle or continue to cut the road to queue for the vehicles in the other lanes. In the second row we see an actor making a right turn in an intersection, where the model correctly predicts that the actor is planning to enter its own lane. However, uncertainty increases compared to the first example, as the vehicle has higher speed as well as heading change rate, and there is a possibility it may enter any of the two vacant lanes. Lastly, in the third row we have a fast actor going straight, while changing lanes to avoid an obstacle in its current lane. The lane change is correctly predicted, as well as lower cross-track uncertainty due to actor's higher speed. Quite intuitively, probability that the actor hits the obstacle is estimated at near-zero.

We conducted deeper analysis of the model behavior in the above three cases. First, we performed dropout analysis to estimate epistemic uncertainty in the model itself [16], by dropping out 50% of randomly selected weights in the fully-connected layers, repeating the process 100 times, and visualizing variance of the resulting trajectory points. The results are shown in the third column




Figure 4: Analysis of the best performing model for three case studies, with results overlaid over input raster images; first column shows ground truth (dotted green line) and predicted trajectory (dotted blue line), second column shows aleatoric uncertainty output by the model, third column shows epistemic uncertainty estimated by dropout analysis, fourth column shows relevant parts of raster estimated by occlusion sensitivity analysis; state inputs are provided above rasters in the first column, indicating velocity (v) in m/s , acceleration (a) in m/s^2 , heading change rate (hcr) in deg/s

of Figure 4, where we see that epistemic uncertainty is very low in all cases, in fact several orders of magnitude lower than aleatoric uncertainty visualized in the second column of Figure 4. This indicates that the model converged, that adding more data would have limited effect on performance, and that the overall uncertainty can be approximated by considering only the learned uncertainty.

Furthermore, we performed occlusion sensitivity analysis to better understand which parts of the raster the model is focusing on during inference [44]. To do so we swept an 15×15 black box across the raster and visualized the amount of change in the output compared to a non-occluded raster (as measured by the average displacement error), with results shown in the fourth column of Figure 4. In the first case the model focuses on the oncoming lane and vehicles in front of the actor, as those parts of the raster are most relevant for a vehicle cutting across oncoming traffic and queuing. Quite intuitively, in the second case the model focuses on nearby vehicles and the crosswalks in the turn lane, while in the third case it focuses on the obstacle and the lane further ahead due to actor's higher speed. Such analysis helps debug and understand what the model learned, and confirms it managed to extract knowledge from the training data that comes natural to experienced human drivers.

5 Conclusion

We presented an efficient and effective solution to a critical part of the AV problem, prediction of future motion of traffic actors. We introduced a deep learning-based methodology that provides both point estimates of future actor positions and their uncertainty, allowing for deeper understanding of traffic state. The proposed method is based on rasterization of each actor's context, followed by training of deep convolutional networks to use the resulting raster images to predict actor's trajectory and its corresponding uncertainty. Extensive evaluation of the methodology strongly suggests its practical benefits, and the framework was subsequently deployed to a fleet of autonomous vehicles.

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015.
- [2] A. Alahi, K. Goel, et al. *Social LSTM: Human Trajectory Prediction in Crowded Spaces*. IEEE, Jun 2016.
- [3] L. Ballan, F. Castaldo, A. Alahi, F. Palmieri, and S. Savarese. *Knowledge Transfer for Scene-Specific Motion Prediction*, page 697–713. Springer International Publishing, 2016.
- [4] L. J. Blincoe, T. R. Miller, E. Zaloshnja, and B. A. Lawrence. The economic and societal impact of motor vehicle crashes, 2010 (revised). Technical Report DOT HS 812 013, National Highway Traffic Safety Administration, May 2015.
- [5] A. Z. Broder, P. Ciccolo, et al. Search advertising using web relevance feedback. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, pages 1013–1022. ACM, 2008.
- [6] S. Chen. Kalman filter for robot vision: a survey. *IEEE Transactions on Industrial Electronics*, 59(11):4409–4420, 2012.
- [7] A. Cosgun, L. Ma, et al. Towards full automated drive in urban environments: A demonstration in gomentum station, california. In *IEEE Intelligent Vehicles Symposium*, pages 1811–1818, 2017.
- [8] L. Fletcher, S. Teller, et al. *The MIT – Cornell Collision and Why It Happened*, page 509–548. Springer Berlin Heidelberg, 2009.
- [9] K. Fragkiadaki, P. Agrawal, et al. Learning visual predictive models of physics for playing billiards. In *International Conference on Learning Representations (ICLR)*, 2016.
- [10] C. Gong and D. McNally. A methodology for automated trajectory prediction analysis. In *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2004.
- [11] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
- [12] Y. N. Harari. Reboot for the ai revolution. *Nature News*, 550(7676):324, 2017.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] N. L. Johnson. The folded normal distribution: Accuracy of estimation by maximum likelihood. *Technometrics*, 4(2):249–256, 1962.
- [15] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D):35–45, 1960.
- [16] A. Kendall and Y. Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems*, 2017.
- [17] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [18] K. M. Kitani, B. D. Ziebart, J. A. Bagnell, and M. Hebert. *Activity Forecasting*, page 201–214. Springer Berlin Heidelberg, 2012.
- [19] J. Kong, M. Pfeiffer, G. Schildbach, and F. Borrelli. Kinematic and dynamic vehicle models for autonomous driving control design. In *Intelligent Vehicles Symposium (IV), 2015 IEEE*, pages 1094–1099. IEEE, 2015.
- [20] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [21] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, 2017.
- [22] N. Lee, W. Choi, P. Vernaza, C. B. Choy, et al. *DESIRE: Distant Future Prediction in Dynamic Scenes with Interacting Agents*. IEEE, Jul 2017.
- [23] S. Lefèvre, D. Vasquez, and C. Laugier. A survey on motion prediction and risk assessment for intelligent vehicles. *ROBOMECH Journal*, 1(1), Jul 2014.

- [24] NCHS. Health, United States, 2016: With chartbook on long-term trends in health. Technical Report 1232, National Center for Health Statistics, May 2017.
- [25] NHTSA. Early estimate of motor vehicle traffic fatalities for the first half (jan–jun) of 2017. Technical Report DOT HS 812 453, National Highway Traffic Safety Administration, December 2017.
- [26] G. Nuti, M. Mirghaemi, P. Treleaven, and C. Yingsaeree. Algorithmic trading. *Computer*, 44(11):61–69, 2011.
- [27] P. Ondruška, J. Dequaire, D. Z. Wang, and I. Posner. End-to-end tracking and semantic segmentation using recurrent neural networks. *arXiv preprint arXiv:1604.05091*, 2016.
- [28] P. Ondruška and I. Posner. Deep tracking: Seeing beyond seeing using recurrent neural networks. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 3361–3367. AAAI Press, 2016.
- [29] S. E. Polzin. Implications to public transportation of emerging technologies. 2016.
- [30] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation. *arXiv preprint arXiv:1801.04381*, 2018.
- [31] M. Schreier, V. Willert, and J. Adamy. An integrated approach to maneuver-based trajectory prediction and criticality assessment in arbitrary road environments. *IEEE Transactions on Intelligent Transportation Systems*, 17(10):2751–2766, Oct 2016.
- [32] A. Sergeev and M. D. Balso. Horovod: fast and easy distributed deep learning in tensorflow. *arXiv preprint arXiv:1802.05799*, 2018.
- [33] D. Silver, A. Huang, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [34] D. Silver, T. Hubert, et al. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*, 2017.
- [35] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [36] S. Singh. Critical reasons for crashes investigated in the national motor vehicle crash causation survey. Technical Report DOT HS 812 115, National Highway Traffic Safety Administration, February 2015.
- [37] T. Streubel and K. H. Hoffmann. *Prediction of driver intended path at intersections*. IEEE, Jun 2014.
- [38] E. J. Topol. *The patient will see you now: the future of medicine is in your hands*. Tantor Media, 2015.
- [39] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):652–663, 2017.
- [40] E. A. Wan and R. Van Der Merwe. The unscented kalman filter for nonlinear estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, pages 153–158. Ieee, 2000.
- [41] J. Wang, D. Fleet, and A. Hertzmann. Gaussian process dynamical models for human motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):283–298, Feb 2008.
- [42] K.-C. Wang and R. Zemel. Classifying nba offensive plays using neural networks. In *Proceedings of MIT Sloan Sports Analytics Conference*, pages 1094–1099, 2016.
- [43] J. Wiest. *Statistical long-term motion prediction*. Universität Ulm, 2017.
- [44] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.
- [45] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey. Maximum entropy inverse reinforcement learning. In *Proceedings of the 23rd national conference on Artificial intelligence-Volume 3*, pages 1433–1438. AAAI Press, 2008.
- [46] J. Ziegler, P. Bender, M. Schreiber, et al. Making bertha drive—an autonomous journey on a historic route. *IEEE Intelligent Transportation Systems Magazine*, 6:8–20, 10 2015.