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

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

DEFENDANT'S NOTICE OF SUBPOENA TO WAYMO LLC

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 Waymo, LLC pursuant to the attached Subpoena. A true and correct copy of the Subpoena is attached as Attachment A.

Perceptive Exhibit 2047
IPR2025-01576

Dated: December 10, 2025

Respectfully submitted,

By: */s/ Roger Fulghum*

Roger Fulghum
Lead Attorney
Texas State Bar No. 00790724
Baker Botts LLP
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 LLP
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.

Perceptive Exhibit 2047
IPR2025-01576

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: Waymo, LLC, 1600 Amphitheater Parkway, Mountain View, CA 94043
(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:	Date and Time:
Baker Botts L.L.P., 1001 Page Mill Road, Building One, Suite 200, Palo Alto, CA 94304-1007	01/20/2026 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

Signature of Clerk or Deputy Clerk

/s/ Roger Fulghum
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).

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.

For access to subpoena materials, see Fed. R. Civ. P. 45(a) Committee Note (2013).

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 20, 2026. Tesla is also seeking deposition testimony from Waymo on the documents and topics set out below. The time and date for the deposition is January 20, 2026 at 9:00 am. Tesla is willing to negotiate with Waymo 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 "Waymo" mean Waymo LLC, 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 "Waymo Prior Art Vehicles" means any vehicle with self-driving or autonomous features that was designed, developed, tested, operated, offered for sale, used in commerce, or otherwise used by Waymo, whether in the United States or another country, before February 6, 2019 including, but not limited to, the Google Lexus RX 450h, the Google Firefly, and the Waymo Chrysler Pacifica Hybrid 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

Perceptive Exhibit 2047
IPR2025-01576

(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.

Perceptive Exhibit 2047
IPR2025-01576

REQUESTS FOR PRODUCTION

1. Documents sufficient to identify all Waymo 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 Waymo Prior Art Vehicle, documents sufficient to establish when such Waymo 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 (including inputs and outputs to) all machine learning models, algorithms, routines, methods, functions, or structures in relation to or within each Waymo Prior Art Vehicle.

4. Documents sufficient to show the operation of ChaufferNet. *See* Attachment C.

5. Documents sufficient to show when ChaufferNet was first in public use, publicly demonstrated, offered for sale, used in commerce, or otherwise available to the public. *See* Attachment C.

6. Documents sufficient to show the implementations of ChaufferNet with Waymo Prior Art Vehicles.

7. Documents sufficient to show how Waymo collected, stored, labeled, and/or generated training data used to train machine learning models implemented in or executed by the Waymo Prior Art Vehicles.

8. Documents sufficient to show the sensors (e.g., cameras, radar, and LiDAR) implemented on Waymo Prior Art Vehicles.

9. Documents sufficient to show how the Waymo Prior Art Vehicles "detect[] and classif[y] objects on the road, while estimating their speed, heading, and acceleration over time." *See* Attachment D at 15.

10. Documents sufficient to show how the Waymo Prior Art Vehicles "model, predict, and understand the intent of each object on the road." *See* Attachment D at 15.

11. Documents sufficient to show how the Waymo Prior Art Vehicles "consider[] all the information [the] software has gathered from perception and behavior prediction, and plot[] out a path for [] vehicles." *See* Attachment D at 15.

12. Documents sufficient to identify all vehicles, products, or services related to U.S. Patent No. 10,496,091, including all such vehicles, products, or services implementing U.S. Patent No. 10,496,091, whether in whole or in part.

DEPOSITION TOPICS

1. For the documents produced in response, testimony concerning the authenticity of each document and whether each document (A) was made at or near the time by — or from information transmitted by — someone with knowledge; (B) was kept in the course of a regularly conducted activity of a business, organization, occupation, or calling, whether or not for profit; and (C) was made as a regular practice of that activity.

2. The identity of all Waymo 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 (including inputs and outputs to) all machine learning models, algorithms, routines, methods, functions, or structures in relation to or within each Waymo Prior Art Vehicles.

Perceptive Exhibit 2047
IPR2025-01576

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

5. The operation and implementation of ChaufferNet in relation to Waymo Prior Art Vehicles. *See* Attachment C

6. The collection, storage, labeling, and/or generating of training data used to train machine learning models executed by or implemented in Waymo Prior Art Vehicles.

7. The sensors (e.g., cameras, radar, and LiDAR) implemented on the Waymo Prior Art Vehicles.

8. Attachment D, including the contents and the circumstances around its authorship and publication.

9. How Waymo Prior Art Vehicles "model, predict, and understand the intent of each object on the road." *See* Attachment D at 15.

10. How Waymo Prior Art Vehicles "model, predict, and understand the intent of each object on the road." *See* Attachment D at 15.

11. How Waymo Prior Art Vehicles "consider[] all the information [the] software has gathered from perception and behavior prediction, and plot[] out a path for [] vehicles." *See* Attachment D at 15.

12. All vehicles, products, or services related to U.S. Patent No. 10,496,091, including all such vehicles, products, or services implementing U.S. Patent No. 10,496,091, 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,	§	
	§	
Plaintiff,	§	
	§	Case No. 2:25-cv-00742-JRG
v.	§	
	§	JURY TRIAL DEMANDED
TESLA, INC.	§	
	§	
Defendant.	§	

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

Perceptive Exhibit 2047
IPR2025-01576

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.

Perceptive Exhibit 2047
IPR2025-01576

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 experts[2] (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.

Perceptive Exhibit 2047
IPR2025-01576

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

Perceptive Exhibit 2047
IPR2025-01576

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

Perceptive Exhibit 2047
IPR2025-01576

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

Perceptive Exhibit 2047
IPR2025-01576

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,

Perceptive Exhibit 2047
IPR2025-01576

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 _____

Perceptive Exhibit 2047
IPR2025-01576

Attachment C

ChauffeurNet: Learning to Drive by Imitating the Best and Synthesizing the Worst

Mayank Bansal MAYBAN@WAYMO.COM
Waymo Research
Mountain View, CA 94043, USA

Alex Krizhevsky* AKRIZHEVSKY@GMAIL.COM

Abhijit Ogale OGALE@WAYMO.COM
Waymo Research
Mountain View, CA 94043, USA

Editor:

arXiv:1812.03079v1 [cs.RO] 7 Dec 2018

Abstract

Our goal is to train a policy for autonomous driving via imitation learning that is robust enough to drive a real vehicle. We find that standard behavior cloning is insufficient for handling complex driving scenarios, even when we leverage a perception system for preprocessing the input and a controller for executing the output on the car: 30 million examples are still not enough. We propose exposing the learner to synthesized data in the form of perturbations to the expert's driving, which creates interesting situations such as collisions and/or going off the road. Rather than purely imitating all data, we augment the imitation loss with additional losses that penalize undesirable events and encourage progress – the perturbations then provide an important signal for these losses and lead to robustness of the learned model. We show that the *ChauffeurNet* model can handle complex situations in simulation, and present ablation experiments that emphasize the importance of each of our proposed changes and show that the model is responding to the appropriate causal factors. Finally, we demonstrate the model driving a car in the real world.

Keywords: Deep Learning, Mid-to-mid Driving, Learning to Drive, Trajectory Prediction.

1. Introduction

In order to drive a car, a driver needs to see and understand the various objects in the environment, predict their possible future behaviors and interactions, and then plan how to control the car in order to safely move closer to their desired destination while obeying the rules of the road. This is a difficult robotics challenge that humans solve well, making imitation learning a promising approach. Our work is about getting imitation learning to the level where it has a shot at driving a real vehicle; although the same insights may apply to other domains, these domains might have different constraints and opportunities, so we do not want to claim contributions there.

∗. Work done while at Google Brain & Waymo.

Perceptive Exhibit 2047
IPR2025-01576

We built our system based on leveraging the training data (30 million real-world expert driving examples, corresponding to about 60 days of continual driving) as effectively as possible. There is a lot of excitement for end-to-end learning approaches to driving which typically focus on learning to directly predict raw control outputs such as steering or braking after consuming raw sensor input such as camera or lidar data. But to reduce sample complexity, we opt for mid-level input and output representations that take advantage of perception and control components. We use a perception system that processes raw sensor information and produces our input: a top-down representation of the environment and intended route, where objects such as vehicles are drawn as oriented 2D boxes along with a rendering of the road information and traffic light states. We present this mid-level input to a recurrent neural network (RNN), named *ChauffeurNet*, which then outputs a driving trajectory that is consumed by a controller which translates it to steering and acceleration. The further advantage of these mid-level representations is that the net can be trained on real or simulated data, and can be easily tested and validated in closed-loop simulations before running on a real car.

Our first finding is that even with 30 million examples, and even with mid-level input and output representations that remove the burden of perception and control, pure imitation learning is not sufficient. As an example, we found that this model would get stuck or collide with another vehicle parked on the side of a narrow street, when a nudging and passing behavior was viable. The key challenge is that we need to run the system closed-loop, where errors accumulate and induce a shift from the training distribution (Ross et al. (2011)). Scientifically, this result is valuable evidence about the limitations of pure imitation in the driving domain, especially in light of recent promising results for high-capacity models (Laskey et al. (2017a)). But practically, we needed ways to address this challenge without exposing demonstrators to new states actively (Ross et al. (2011); Laskey et al. (2017b)) or performing reinforcement learning (Kuefler et al. (2017)).

We find that this challenge is surmountable if we augment the imitation loss with losses that discourage bad behavior and encourage progress, and, importantly, augment our data with *synthesized* perturbations in the driving trajectory. These expose the model to non-expert behavior such as collisions and off-road driving, and inform the added losses, teaching the model to avoid these behaviors. Note that the opportunity to synthesize this data comes from the mid-level input-output representations, as perturbations would be difficult to generate with either raw sensor input or direct controller outputs.

We evaluate our system, as well as the relative importance of both loss augmentation and data augmentation, first in simulation. We then show how our final model successfully drives a car in the real world and is able to negotiate situations involving other agents, turns, stop signs, and traffic lights. Finally, it is important to note that there are highly interactive situations such as merging which may require a significant degree of exploration within a reinforcement learning (RL) framework. This will demand simulating other (human) traffic participants, a rich area of ongoing research. Our contribution can be viewed as pushing the boundaries of what you can do with purely offline data and no RL.

Perceptive Exhibit 2047
IPR2025-01576

2. Related Work

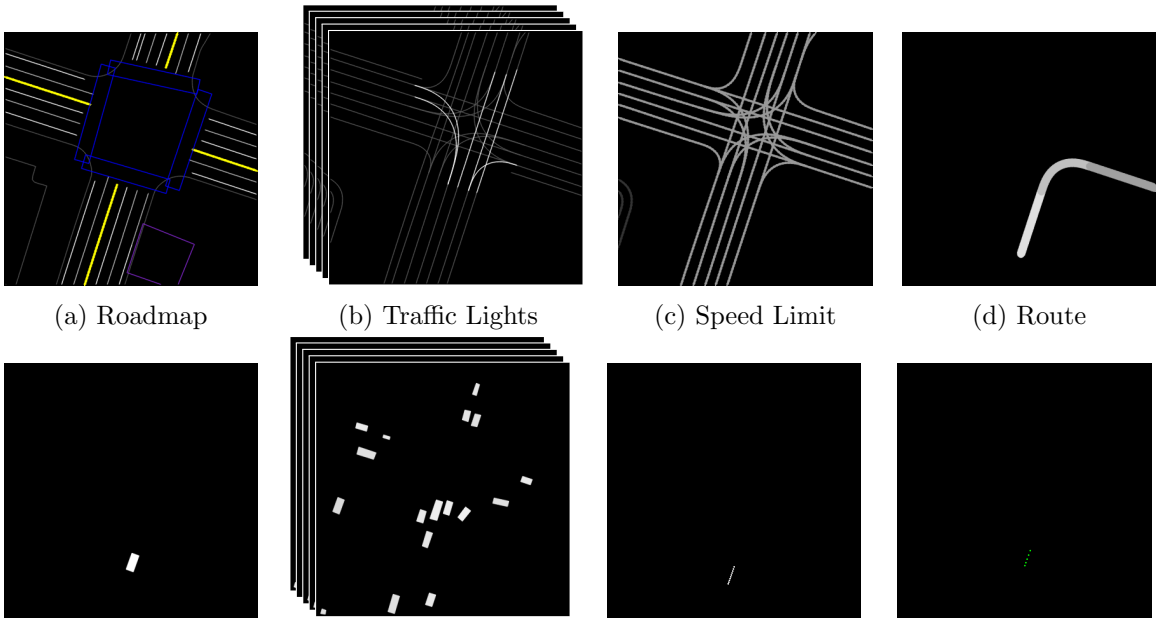
Decades-old work on ALVINN (Pomerleau (1989)) showed how a shallow neural network could follow the road by directly consuming camera and laser range data. Learning to drive in an end-to-end manner has seen a resurgence in recent years. Recent work by Chen et al. (2015) demonstrated a convolutional net to estimate affordances such as distance to the preceding car that could be used to program a controller to control the car on the highway. Researchers at NVIDIA (Bojarski et al. (2016, 2017)) showed how to train an end-to-end deep convolutional neural network that steers a car by consuming camera input. Xu et al. (2017) trained a neural network for predicting discrete or continuous actions also based on camera inputs. Codevilla et al. (2018) also train a network using camera inputs and conditioned on high-level commands to output steering and acceleration. Kuefler et al. (2017) use Generative Adversarial Imitation Learning (GAIL) with simple affordance-style features as inputs to overcome cascading errors typically present in behavior cloned policies so that they are more robust to perturbations. Recent work from Hecker et al. (2018) learns a driving model using 360-degree camera inputs and desired route planner to predict steering and speed. The CARLA simulator (Dosovitskiy et al. (2017)) has enabled recent work such as Sauer et al. (2018), which estimates several affordances from sensor inputs to drive a car in a simulated urban environment. Using mid-level representations in a spirit similar to our own, Müller et al. (2018) train a system in simulation using CARLA by training a driving policy from a scene segmentation network to output high-level control, thereby enabling transfer learning to the real world using a different segmentation network trained on real data. Pan et al. (2017) also describes achieving transfer of an agent trained in simulation to the real world using a learned intermediate scene labeling representation. Reinforcement learning may also be used in a simulator to train drivers on difficult interactive tasks such as merging which require a lot of exploration, as shown in Shalev-Shwartz et al. (2016). A convolutional network operating on a space-time volume of bird's eye-view representations is also employed by Luo et al. (2018); Djuric et al. (2018); Lee et al. (2017) for tasks like 3D detection, tracking and motion forecasting. Finally, there exists a large volume of work on vehicle motion planning outside the machine learning context and Paden et al. (2016) present a notable survey.

3. Model Architecture

3.1 Input Output Representation

We begin by describing our top-down input representation that the network will process to output a drivable trajectory. At any time t, our agent (or vehicle) may be represented in a top-down coordinate system by $\mathbf{p}_t, \theta_t, s_t$, where $\mathbf{p}_t = (x_t, y_t)$ denotes the agent's location or pose, θ_t denotes the heading or orientation, and s_t denotes the speed. The top-down coordinate system is picked such that our agent's pose $\mathbf{p}_0$ at the current time $t = 0$ is always at a fixed location (u_0, v_0) within the image. For data augmentation purposes during training, the orientation of the coordinate system is randomly picked for each training example to be within an angular range of $\theta_0 \pm \Delta$, where θ_0 denotes the heading or orientation of our agent at time $t = 0$. The top-down view is represented by a set of images of size $W \times H$ pixels, at a ground sampling resolution of ϕ meters/pixel. Note that as the agent

Perceptive Exhibit 2047
IPR2025-01576



(a) Roadmap (b) Traffic Lights (c) Speed Limit (d) Route

(e) Current Agent Box (f) Dynamic Boxes (g) Past Agent Poses (h) Future Agent Poses

Figure 1: Driving model inputs (a-g) and output (h).

moves, this view of the environment moves with it so the agent always sees a fixed forward range, $R_{forward} = (H - v_0)\phi$ of the world – similar to having an agent with sensors that see only up to $R_{forward}$ meters forward.

As shown in Fig. 1, the input to our model consists of several images of size $W \times H$ pixels rendered into this top-down coordinate system. (a) Roadmap: a color (3-channel) image with a rendering of various map features such as lanes, stop signs, cross-walks, curbs, etc. (b) Traffic lights: a temporal sequence of grayscale images where each frame of the sequence represents the known state of the traffic lights at each past timestep. Within each frame, we color each lane center by a gray level with the brightest level for red lights, intermediate gray level for yellow lights, and a darker level for green or unknown lights[1]. (c) Speed limit: a single channel image with lane centers colored in proportion to their known speed limit. (d) Route: the intended route along which we wish to drive, generated by a router (think of a Google Maps-style route). (e) Current agent box: this shows our agent's full bounding box at the current timestep $t = 0$. (f) Dynamic objects in the environment: a temporal sequence of images showing all the potential dynamic objects (vehicles, cyclists, pedestrians) rendered as oriented boxes. (g) Past agent poses: the past poses of our agent are rendered into a single grayscale image as a trail of points.

We use a fixed-time sampling of δt to sample any past or future temporal information, such as the traffic light state or dynamic object states in the above inputs. The traffic

1. We employ an indexed representation for roadmap and traffic lights channels to reduce the number of input channels, and to allow extensibility of the input representation to express more roadmap features or more traffic light states without changing the model architecture.

Perceptive Exhibit 2047
IPR2025-01576

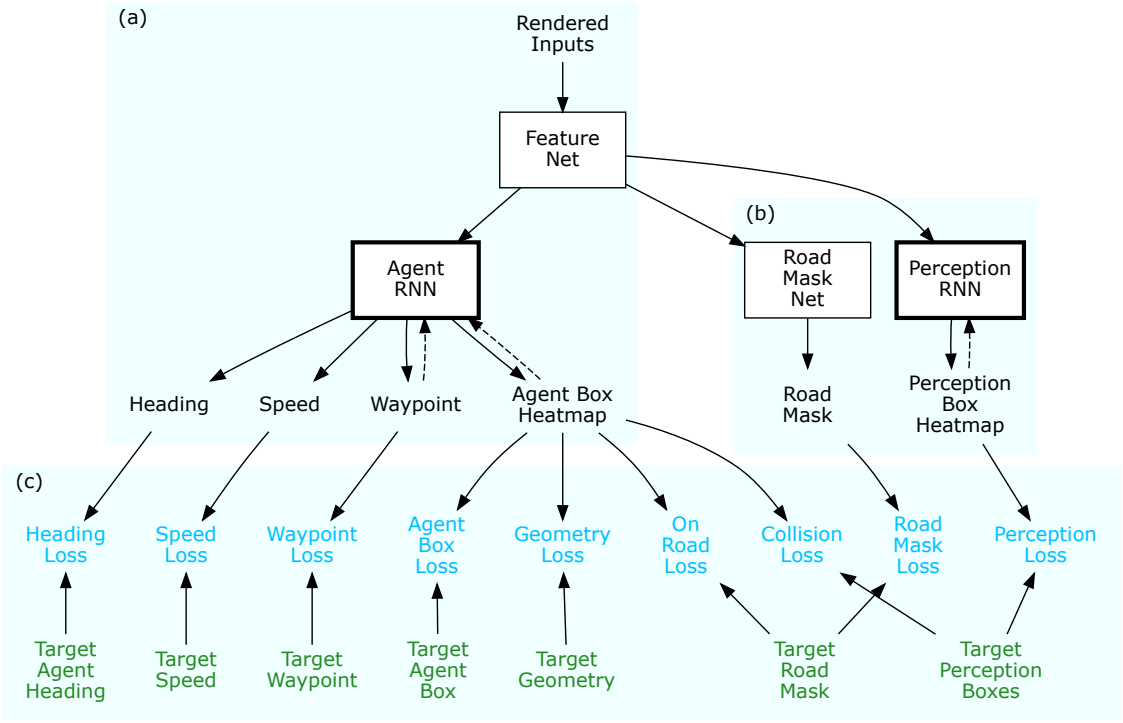


Figure 2: Training the driving model. (a) The core *ChauffeurNet* model with a FeatureNet and an AgentRNN, (b) Co-trained road mask prediction net and PerceptionRNN, and (c) Training losses are shown in blue, and the green labels depict the ground-truth data. The dashed arrows represent the recurrent feedback of predictions from one iteration to the next.

lights and dynamic objects are sampled over the past T_{scene} seconds, while the past agent poses are sampled over a potentially longer interval of T_{pose} seconds. This simple input representation, particularly the box representation of other dynamic objects, makes it easy to generate input data from simulation or create it from real-sensor logs using a standard perception system that detects and tracks objects. This enables testing and validation of models in closed-loop simulations before running them on a real car. This also allows the same model to be improved using simulated data to adequately explore rare situations such as collisions for which real-world data might be difficult to obtain. Using a top-down 2D view also means efficient convolutional inputs, and allows flexibility to represent metadata and spatial relationships in a human-readable format. Papers on testing frameworks such as Tian et al. (2018), Pei et al. (2017) show the brittleness of using raw sensor data (such as camera images or lidar point clouds) for learning to drive, and reinforce the approach of using an intermediate input representation.

If I denotes the set of all the inputs enumerated above, then the *ChauffeurNet* model recurrently predicts future poses of our agent conditioned on these input images I as shown by the green dots in Fig. 1(h).

$$\mathbf{p}_{t+\delta t} = \text{ChauffeurNet}(I, \mathbf{p}_t) \tag{1}$$

Perceptive Exhibit 2047
IPR2025-01576

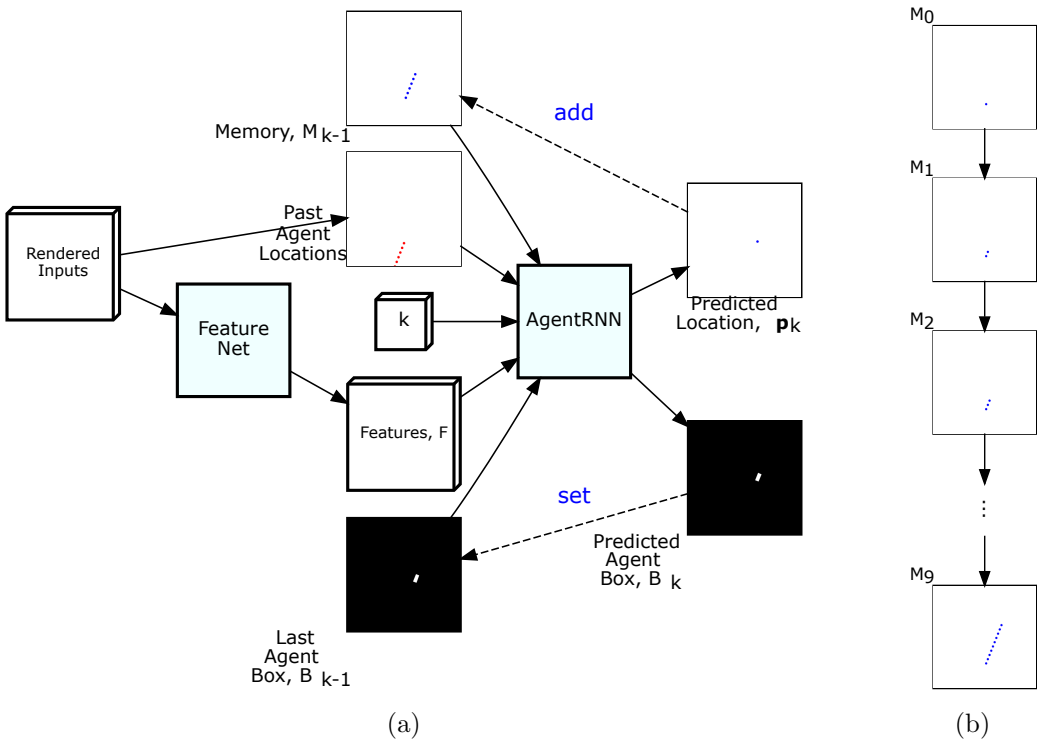


Figure 3: (a) Schematic of ChauffeurNet. (b) Memory updates over multiple iterations.

In Eq. (1), current pose $\mathbf{p}_0$ is a known part of the input, and then the ChauffeurNet performs N iterations and outputs a future trajectory$\{\mathbf{p}_{\delta t}, \mathbf{p}_{2\delta t}, ..., \mathbf{p}_{N\delta t}\}$ along with other properties such as future speeds. This trajectory can be fed to a controls optimizer that computes detailed driving control (such as steering and braking commands) within the specific constraints imposed by the dynamics of the vehicle to be driven. Different types of vehicles may possibly utilize different control outputs to achieve the same driving trajectory, which argues against training a network to directly output low-level steering and acceleration control. Note, however, that having intermediate representations like ours does not preclude end-to-end optimization from sensors to controls.

3.2 Model Design

Broadly, the driving model is composed of several parts as shown in Fig. 2. The main *ChauffeurNet* model shown in part (a) of the figure consists of a convolutional feature network (*FeatureNet*) that consumes the input data to create a digested contextual feature representation that is shared by the other networks. These features are consumed by a recurrent agent network (*AgentRNN*) that iteratively predicts successive points in the driving trajectory. Each point at time t in the trajectory is characterized by its location $\mathbf{p}_t = (x_t, y_t)$, heading θ_t and speed s_t. The *AgentRNN* also predicts the bounding box of the vehicle as a spatial heatmap at each future timestep. In part (b) of the figure, we see that two other networks are co-trained using the same feature representation as an input. The Road Mask Network predicts the drivable areas of the field of view (on-road vs. off-road), while the

Perceptive Exhibit 2047
IPR2025-01576

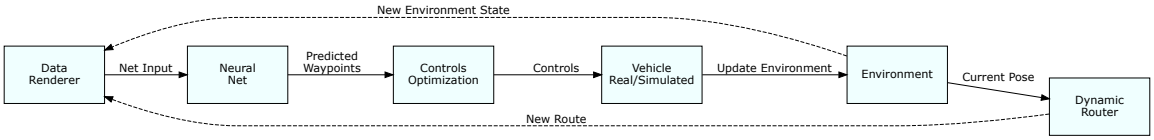


Figure 4: Software architecture for the end-to-end driving pipeline.

recurrent perception network (*PerceptionRNN*) iteratively predicts a spatial heatmap for each timestep showing the future location of every other agent in the scene. We believe that doing well on these additional tasks using the same shared features as the main task improves generalization on the main task. Fig. 2(c) shows the various losses used in training the model, which we will discuss in detail below.

Fig. 3 illustrates the ChauffeurNet model in more detail. The rendered inputs shown in Fig. 1 are fed to a large-receptive field convolutional *FeatureNet* with skip connections, which outputs features F that capture the environmental context and the intent. These features are fed to the *AgentRNN* which predicts the next point $\mathbf{p}_k$ on the driving trajectory, and the agent bounding box heatmap B_k, conditioned on the features F from the FeatureNet, the iteration number $k \in \{1, \ldots, N\}$, the memory M_{k-1} of past predictions from the *AgentRNN*, and the agent bounding box heatmap B_{k-1} predicted in the previous iteration.

$$\mathbf{p}_k, B_k = \text{AgentRNN}(k, F, M_{k-1}, B_{k-1}) \tag{2}$$

The memory M_k is an additive memory consisting of a single channel image. At iteration k of the *AgentRNN*, the memory is incremented by 1 at the location $\mathbf{p}_k$ predicted by the *AgentRNN*, and this memory is then fed to the next iteration. The *AgentRNN* outputs a heatmap image over the next pose of the agent, and we use the arg-max operation to obtain the coarse pose prediction $\mathbf{p}_k$ from this heatmap. The *AgentRNN* then employs a shallow convolutional meta-prediction network with a fully-connected layer that predicts a sub-pixel refinement of the pose $\delta\mathbf{p}_k$ and also estimates the heading θ_k and the speed s_k. Note that the *AgentRNN* is unrolled at training time for a fixed number of iterations, and the losses described below are summed together over the unrolled iterations. This is possible because of the non-traditional RNN design where we employ an explicitly crafted memory model instead of a learned memory.

3.3 System Architecture

Fig. 4 shows a system level overview of how the neural net is used within the self-driving system. At each time, the updated state of our agent and the environment is obtained via a perception system that processes sensory output from the real-world or from a simulation environment as the case may be. The intended route is obtained from the router, and is updated dynamically conditioned on whether our agent was able to execute past intents or not. The environment information is rendered into the input images described in Fig. 1 and given to the RNN which then outputs a future trajectory. This is fed to a controls optimizer that outputs the low-level control signals that drive the vehicle (in the real world or in simulation).

Perceptive Exhibit 2047
IPR2025-01576

4. Imitating the Expert

In this section, we first show how to train the model above to imitate the expert.

4.1 Imitation Losses

4.1.1 Agent Position, Heading and Box Prediction

The *AgentRNN* produces three outputs at each iteration k: a probability distribution $P_k(x, y)$ over the spatial coordinates of the predicted waypoint obtained after a spatial *softmax*, a heatmap of the predicted agent box at that timestep $B_k(x, y)$ obtained after a per-pixel *sigmoid* activation that represents the probability that the agent occupies a particular pixel, and a regressed box heading output θ_k. Given ground-truth data for the above predicted quantities, we can define the corresponding losses for each iteration as:

$$\mathcal{L}_p = \mathcal{H}(P_k, P_k^{gt}) \tag{3}$$

$$\mathcal{L}_B = \frac{1}{WH} \sum_x \sum_y \mathcal{H}(B_k(x, y), B_k^{gt}(x, y)) \tag{4}$$

$$\mathcal{L}_\theta = \left\| \theta_k - \theta_k^{gt} \right\|_1 \tag{5}$$

where the superscript gt denotes the corresponding ground-truth values, and $\mathcal{H}(a, b)$ is the cross-entropy function. Note that P_k^{gt} is a binary image with only the pixel at the ground-truth target coordinate $\lfloor \mathbf{p}_k^{gt} \rfloor$ set to one.

4.1.2 Agent Meta Prediction

The meta prediction network performs regression on the features to generate a sub-pixel refinement $\delta\mathbf{p}_k$ of the coarse waypoint prediction as well as a speed estimate s_k at each iteration. We employ L_1 loss for both of these outputs:

$$\mathcal{L}_{p-subpixel} = \left\| \delta\mathbf{p}_k - \delta\mathbf{p}_k^{gt} \right\|_1 \tag{6}$$

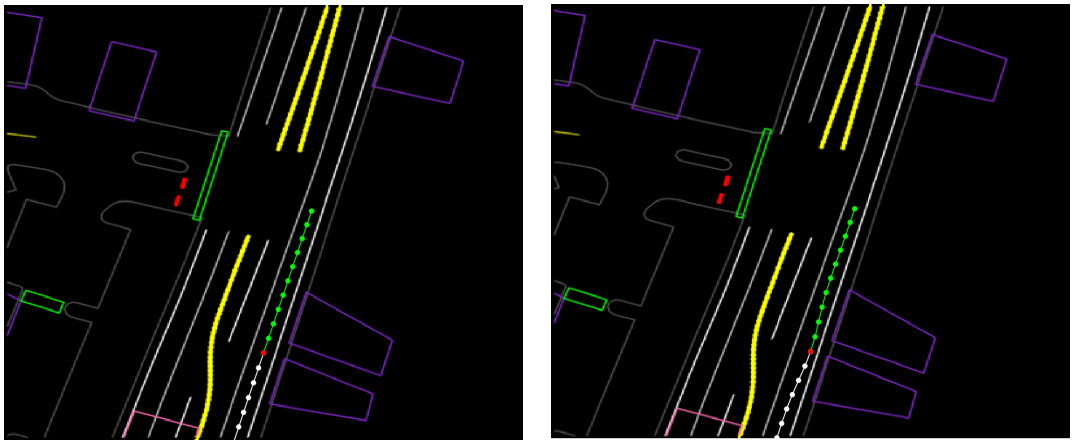
$$\mathcal{L}_{speed} = \left\| s_k - s_k^{gt} \right\|_1 \tag{7}$$

where $\delta\mathbf{p}_k^{gt} = \mathbf{p}_k^{gt} - \lfloor \mathbf{p}_k^{gt} \rfloor$ is the fractional part of the ground-truth pose coordinates.

4.2 Past Motion Dropout

During training, the model is provided the past motion history as one of the inputs (Fig. 1(g)). Since the past motion history during training is from an expert demonstration, the net can learn to "cheat" by just extrapolating from the past rather than finding the underlying causes of the behavior. During closed-loop inference, this breaks down because the past history is from the net's own past predictions. For example, such a trained net may learn to only stop for a stop sign if it sees a deceleration in the past history, and will therefore never stop for a stop sign during closed-loop inference. To address this, we introduce a dropout on the past pose history, where for 50% of the examples, we keep only the current position (u_0, v_0) of the agent in the past agent poses channel of the input data. This forces the net

Perceptive Exhibit 2047
IPR2025-01576



(a) Original (b) Perturbed

Figure 5: Trajectory Perturbation. (a) An original logged training example where the agent is driving along the center of the lane. (b) The perturbed example created by perturbing the current agent location (red point) in the original example away from the lane center and then fitting a new smooth trajectory that brings the agent back to the original target location along the lane center.

to look at other cues in the environment to explain the future motion profile in the training example.

5. Beyond Pure Imitation

In this section, we go beyond vanilla cloning of the expert's demonstrations in order to teach the model to arrest drift and avoid bad behavior such as collisions and off-road driving by synthesizing variations of the expert's behavior.

5.1 Synthesizing Perturbations

Running the model as a part of a closed-loop system over time can cause the input data to deviate from the training distribution. To prevent this, we train the model by adding some examples with realistic perturbations to the agent trajectories. The start and end of a trajectory are kept constant, while a perturbation is applied around the midpoint and smoothed across the other points. Quantitatively, we jitter the midpoint pose of the agent uniformly at random in the range $[-0.5, 0.5]$ meters in both axes, and perturb the heading by $[-\pi/3, \pi/3]$ radians. We then fit a smooth trajectory to the perturbed point and the original start and end points. Such training examples bring the car back to its original trajectory after a perturbation. Fig. 5 shows an example of perturbing the current agent location (red point) away from the lane center and the fitted trajectory correctly bringing it back to the original target location along the lane center. We filter out some perturbed trajectories that are impractical by thresholding on maximum curvature. But we do allow the perturbed trajectories to collide with other agents or drive off-road, because the network can then experience and avoid such behaviors even though real examples of these cases are

Perceptive Exhibit 2047
IPR2025-01576

not present in the training data. In training, we give perturbed examples a weight of 1/10 relative to the real examples, to avoid learning a propensity for perturbed driving.

5.2 Beyond the Imitation Loss

5.2.1 Collision Loss

Since our training data does not have any real collisions, the idea of avoiding collisions is implicit and will not generalize well. To alleviate this issue, we add a specialized loss that directly measures the overlap of the predicted agent box B_k with the ground-truth boxes of all the scene objects at each timestep.

$$\mathcal{L}_{collision} = \frac{1}{WH} \sum_x \sum_y B_k(x,y) \; . \; Obj_k^{gt}(x,y) \tag{8}$$

where B_k is the likelihood map for the output agent box prediction, and Obj_k^{gt} is a binary mask with ones at all pixels occupied by other dynamic objects (other vehicles, pedestrians, etc.) in the scene at timestep k. At any time during training, if the model makes a poor prediction that leads to a collision, the overlap loss would influence the gradients to correct the mistake. However, this loss would be effective only during the initial training rounds when the model hasn't learned to predict close to the ground-truth locations due to the absence of real collisions in the ground truth data. This issue is alleviated by the addition of trajectory perturbation data, where artificial collisions within those examples allow this loss to be effective throughout training without the need for online exploration like in reinforcement learning settings.

5.2.2 On Road Loss

Trajectory perturbations also create synthetic cases where the car veers off the road or climbs a curb or median because of the perturbation. To train the network to avoid hitting such hard road edges, we add a specialized loss that measures overlap of the predicted agent box B_k in each timestep with a binary mask $Road^{gt}$ denoting the road and non-road regions within the field-of-view.

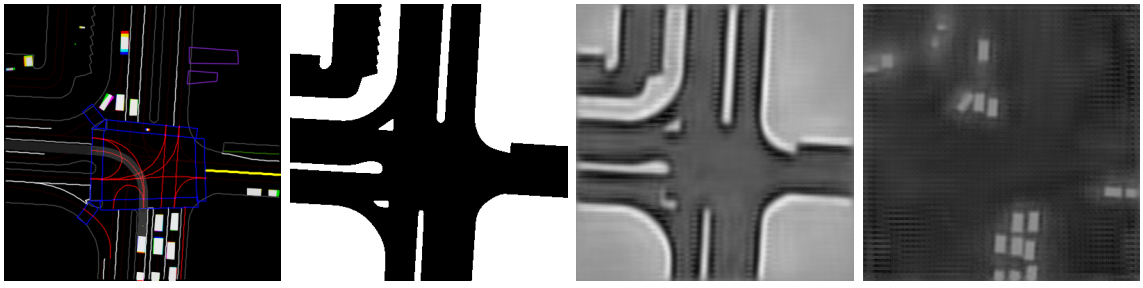
$$\mathcal{L}_{onroad} = \frac{1}{WH} \sum_x \sum_y B_k(x,y) \; . \; (1 - Road^{gt}(x,y)) \tag{9}$$

5.2.3 Geometry Loss

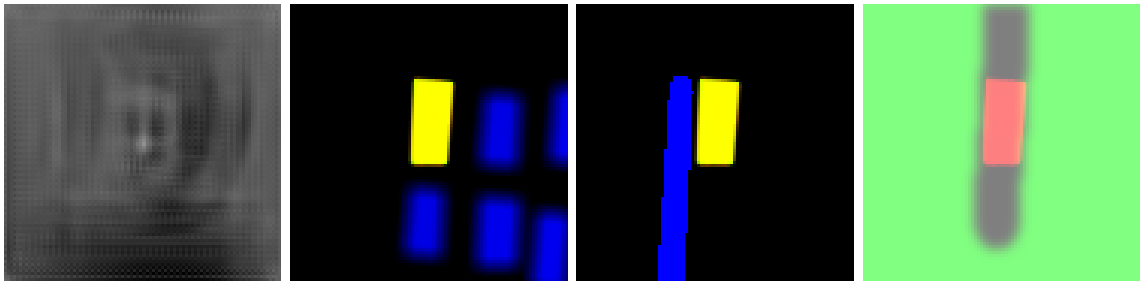
We would like to explicitly constrain the agent to follow the target geometry independent of the speed profile. We model this target geometry by fitting a smooth curve to the target waypoints and rendering this curve as a binary image in the top-down coordinate system. The thickness of this curve is set to be equal to the width of the agent. We express this loss similar to the collision loss by measuring the overlap of the predicted agent box with the binary target geometry image $Geom^{gt}$. Any portion of the box that does not overlap with the target geometry curve is added as a penalty to the loss function.

$$\mathcal{L}_{geom} = \frac{1}{WH} \sum_x \sum_y B_k(x,y) \; . \; (1 - Geom^{gt}(x,y)) \tag{10}$$

Perceptive Exhibit 2047
IPR2025-01576



(a) Flattened Inputs (b) Target Road Mask (c) Pred Road Mask Logits (d) Pred Vehicles Logits



(e) Agent Pose Logits (f) Collision Loss (g) On Road Loss (h) Geometry Loss

Figure 6: Visualization of predictions and loss functions on an example input. The top row is at the input resolution, while the bottom row shows a zoomed-in view around the current agent location.

5.2.4 Auxiliary Losses

Similar to our own agent's trajectory, the motion of other agents may also be predicted by a recurrent network. Correspondingly, we add a recurrent perception network *PerceptionRNN* that uses as input the shared features F created by the *FeatureNet* and its own predictions Obj_{k-1} from the previous iteration, and predicts a heatmap Obj_k at each iteration. $Obj_k(x, y)$ denotes the probability that location (x, y) is occupied by a dynamic object at time k. For iteration $k = 0$, the PerceptionRNN is fed the ground truth objects at the current time.

$$\mathcal{L}_{objects} = \frac{1}{WH} \sum_x \sum_y \mathcal{H}(Obj_k(x, y), Obj_k^{gt}(x, y)) \tag{11}$$

Co-training a *PerceptionRNN* to predict the future of other agents by sharing the same feature representation F used by the *PerceptionRNN* is likely to induce the feature network to learn better features that are suited to both tasks. Several examples of predicted trajectories from *PerceptionRNN* on logged data are shown on our website here.

We also co-train to predict a binary road/non-road mask by adding a small network of convolutional layers to the output of the feature net F. We add a cross-entropy loss to the predicted road mask output $Road(x, y)$ which compares it to the ground-truth road mask

Perceptive Exhibit 2047
IPR2025-01576

T_{scene}	T_{pose}	δt	N	Δ	W	H	u_0	v_0	ϕ
1.0 s	8.0 s	0.2s	10	25°	400 px	400 px	200 px	320 px	0.2 m/px

Table 1: Parameter values for the experiments in this paper.

Rendering	FeatureNet	AgentRNN (N=10)	PerceptionRNN (N=10)	**Overall**
8 ms	6.5 ms	145 ms	35 ms	**160 ms**

Table 2: Run-time performance on NVIDIA Tesla P100 GPU.

$Road^{gt}$.

$$\mathcal{L}_{road} = \frac{1}{WH} \sum_x \sum_y \mathcal{H}(Road(x,y), Road^{gt}(x,y)) \tag{12}$$

Fig. 6 shows some of the predictions and losses for a single example processed through the model.

5.3 Imitation Dropout

Overall, our losses may be grouped into two sub-groups, the imitation losses:

$$\mathcal{L}_{imit} = \{\mathcal{L}_p, \mathcal{L}_B, \mathcal{L}_\theta, \mathcal{L}_{p-subpixel}, \mathcal{L}_{speed}\} \tag{13}$$

and the environment losses:

$$\mathcal{L}_{env} = \{\mathcal{L}_{collision}, \mathcal{L}_{onroad}, \mathcal{L}_{geom}, \mathcal{L}_{objects}, \mathcal{L}_{road}\} \tag{14}$$

The imitation losses cause the model to imitate the expert's demonstrations, while the environment losses discourage undesirable behavior such as collisions. To further increase the effectiveness of the environment losses, we experimented with randomly dropping out the imitation losses for a random subset of training examples. We refer to this as "imitation dropout". In the experiments, we show that imitation dropout yields a better driving model than simply under-weighting the imitation losses. During imitation dropout, the weight on the imitation losses w_{imit} is randomly chosen to be either 0 or 1 with a certain probability for each training example. The overall loss is given by:

$$\mathcal{L} = w_{imit} \sum_{\ell \in \mathcal{L}_{imit}} \ell + w_{env} \sum_{\ell \in \mathcal{L}_{env}} \ell \tag{15}$$

6. Experiments

6.1 Data

The training data to train our model was obtained by randomly sampling segments of real-world expert driving and removing segments where the car was stationary for long periods of time. Our input field of view is $80m \times 80m$ ($W\phi = 80$) and with the agent positioned at (u_0, v_0), we get an effective forward sensing range of $R_{forward} = 64m$. Therefore, for the

Perceptive Exhibit 2047
IPR2025-01576

experiments in this work we also removed any segments of highway driving given the longer sensing range requirement that entails. Our dataset contains approximately 26 million examples which amount to about 60 days of continuous driving. As discussed in Section 3, the vertical-axis of the top-down coordinate system for each training example is randomly oriented within a range of $\Delta = \pm 25^\circ$ of our agent's current heading, in order to avoid a bias for driving along the vertical axis. The rendering orientation is set to the agent heading ($\Delta = 0$) during inference. Data about the prior map of the environment (roadmap) and the speed-limits along the lanes is collected apriori. For the dynamic scene entities like objects and traffic-lights, we employ a separate perception system based on laser and camera data similar to existing works in the literature (Yang et al. (2018); Fairfield and Urmson (2011)). Table 1 lists the parameter values used for all the experiments in this paper. The model runs on a NVidia Tesla P100 GPU in 160ms with the detailed breakdown in Table 2.

6.2 Models

We train and test not only our final model, but a sequence of models that introduce the ingredients we describe one by one on top of behavior cloning. We start with $\mathcal{M}_0$, which does behavior cloning with past motion dropout to prevent using the history to cheat. $\mathcal{M}_1$ adds perturbations without modifying the losses. $\mathcal{M}_2$ further adds our environment losses $\mathcal{L}_{env}$ in Section 5.2. $\mathcal{M}_3$ and $\mathcal{M}_4$ address the fact that we do not want to imitate bad behavior – $\mathcal{M}_3$ is a baseline approach, where we simply decrease the weight on the imitation loss, while $\mathcal{M}_4$ uses our imitation dropout approach with a dropout probability of 0.5. Table 3 lists the configuration for each of these models.

6.3 Closed Loop Evaluation

To evaluate our learned model on a specific scenario, we replay the segment through the simulation until a buffer period of $\max(T_{pose}, T_{scene})$ has passed. This allows us to generate the first rendered snapshot of the model input using all the replayed messages until now. The model is evaluated on this input, and the fitted controls are passed to the vehicle simulator that emulates the dynamics of the vehicle thus moving the simulated agent to its next pose. At this point, the simulated pose might be different from the logged pose, but our input representation allows us to correctly render the new input for the model relative to the new pose. This process is repeated until the end of the segment, and we evaluate scenario specific metrics like stopping for a stop-sign, collision with another vehicle etc. during the simulation. Since the model is being used to drive the agent forward, this is a *closed-loop* evaluation setup.

6.3.1 Model Ablation Tests

Here, we present results from experiments using the various models in the closed-loop simulation setup. We first evaluated all the models on simple situations such as stopping for stop-signs and red traffic lights, and lane following along straight and curved roads by creating 20 scenarios for each situation, and found that *all the models worked well in these simple cases*. Therefore, we will focus below on specific complex situations that highlight the differences between these models.

Perceptive Exhibit 2047
IPR2025-01576

Model	Description	w_{imit}	w_{env}
$\mathcal{M}_0$	Imitation with Past Dropout	1.0	0.0
$\mathcal{M}_1$	$\mathcal{M}_0$ + Traj Perturbation	1.0	0.0
$\mathcal{M}_2$	$\mathcal{M}_1$ + Environment Losses	1.0	1.0
$\mathcal{M}_3$	$\mathcal{M}_2$ with less imitation	0.5	1.0
$\mathcal{M}_4$	$\mathcal{M}_2$ with Imitation Dropout	*Dropout probability = 0.5 (see Section 5.3).*	

Table 3: Model configuration for the model ablation tests.

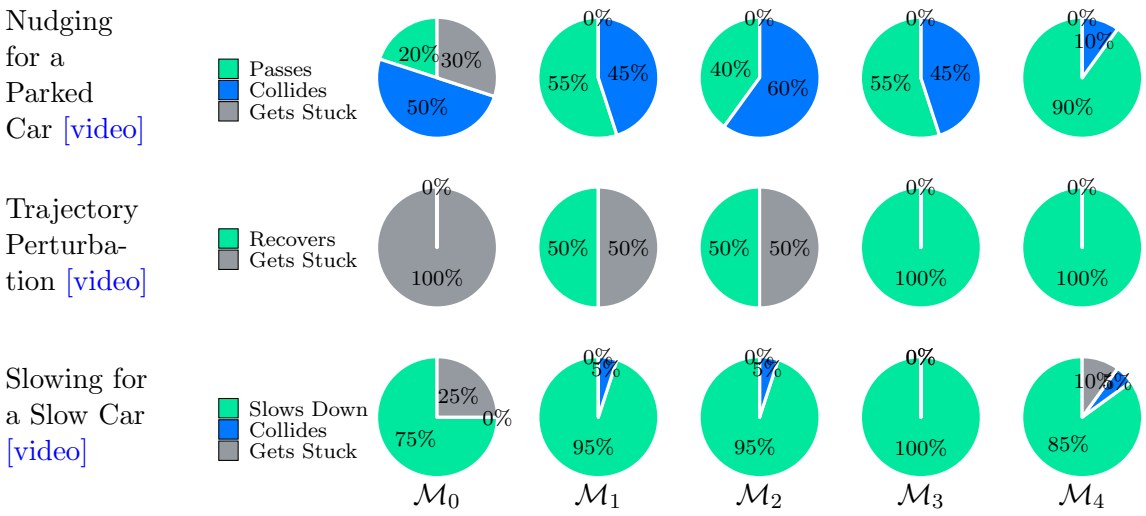


Figure 7: Model ablation test results on three scenario types.

Nudging around a parked car. To set up this scenario, we place the agent at an arbitrary distance from a stop-sign on an undivided two-way street and then place a parked vehicle on the right shoulder between the the agent and the stop-sign. We pick 4 separate locations with both straight and curved roads then vary the starting speed of the agent between 5 different values to create a total of 20 scenarios. We then observe if the agent would stop and get stuck behind, collide with the parked car, or correctly pass around the parked car, and report the aggregate performance in Fig. 7(row 1). We find that other than $\mathcal{M}_4$, all other models cause the agent to collide with the parked vehicle about half the time. The baseline $\mathcal{M}_0$ model can also get stuck behind the parked vehicle in some of the scenarios. The model $\mathcal{M}_4$ nudges around the parked vehicle and then brings the agent back to the lane center. This can be attributed to the model's ability to learn to avoid collisions and nudge around objects because of training with the collision loss the trajectory perturbation. Comparing model $\mathcal{M}_3$ and $\mathcal{M}_4$, it is apparent that "imitation dropout" was more effective at learning the right behavior than only re-weighting the imitation losses. Note that in this scenario, we generate several variations by changing the starting speed of the agent relative to the parked car. This creates situations of increasing difficulty, where the agent approaches the parked car at very high relative speed and thus does not have enough time to nudge around the car given the dynamic constraints. A 10% collision rate for $\mathcal{M}_4$ is thus not a measure of the absolute performance of the model since we do not have a perfect

Perceptive Exhibit 2047
IPR2025-01576

driver which could have performed well at all the scenarios here. But in relative terms, this model performs the best.

Recovering from a trajectory perturbation. To set up this scenario, we place the agent approaching a curved road and vary the starting position and the starting speed of the agent to generate a total of 20 scenario variations. Each variation puts the agent at a different amount of offset from the lane center with a different heading error relative to the lane. We then measure how well the various models are at recovering from the lane departure. Fig. 7(row 2) presents the results aggregated across these scenarios and shows the contrast between the baseline model $\mathcal{M}_0$ which is not able to recover in any of the situations and the models $\mathcal{M}_3$ and $\mathcal{M}_4$ which handle all deviations well. All models trained with the perturbation data are able to handle 50% of the scenarios which have a lower starting speed. At a higher starting speed, we believe that $\mathcal{M}_3$ and $\mathcal{M}_4$ do better than $\mathcal{M}_1$ and $\mathcal{M}_2$ because they place a higher emphasis on the imagination losses.

Slowing down for a slow car. To set up this scenario, we place the agent on a straight road at varying initial speeds and place another car ahead with a varying but slower constant speed, generating a total of 20 scenario variations, to evaluate the ability to slow for and then follow the car ahead. From Fig. 7(row 3), we see that some models slow down to zero speed and get stuck. For the variation with the largest relative speed, there isn't enough time for most models to stop the agent in time, thus leading to a collision. For these cases, model $\mathcal{M}_3$ which uses imitation loss re-weighting works better than the model $\mathcal{M}_4$ which uses imitation dropout. $\mathcal{M}_4$ has trouble in two situations due to being over aggressive in trying to maneuver around the slow car and then grazes the left edge of the road. This happens in the two extreme variations where the relative speed between the two cars is the highest.

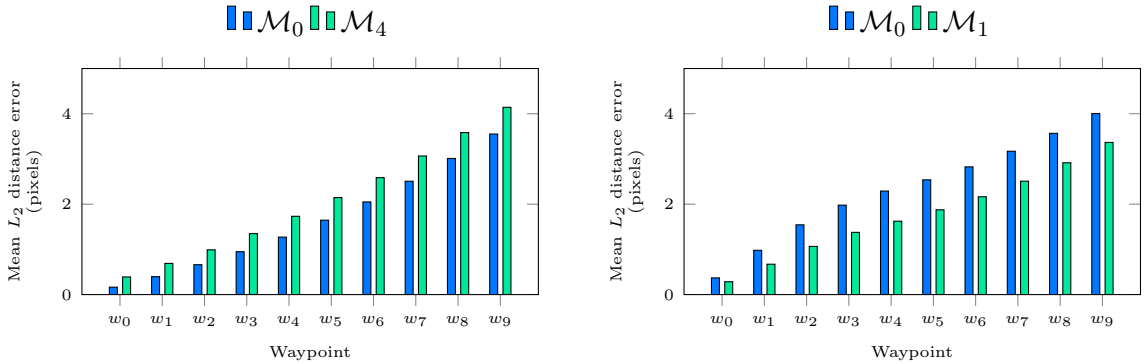
6.3.2 Input Ablation Tests

With input ablation tests, we want to test the final $\mathcal{M}_4$ model's ability to identify the correct causal factors behind specific behaviors, by testing the model's behavior in the presence or absence of the correct causal factor while holding other conditions constant. In simulation, we have evaluated our model on 20 scenarios with and without stop-signs rendered, and 20 scenarios with and without other vehicles in the scene rendered. The model exhibits the correct behavior in all scenarios, thus confirming that it has learned to respond to the correct features for a stop-sign and a stopped vehicle.

6.3.3 Logged Data Simulated Driving

For this evaluation, we take logs from our real-driving test data (separate from our training data), and use our trained network to drive the car using the vehicle simulator keeping everything else the same i.e. the dynamic objects, traffic-light states etc. are all kept the same as in the logs. Some example videos are shown here and they illustrate the ability of the model in dealing with multiple dynamic objects and road controls.

Perceptive Exhibit 2047
IPR2025-01576



(a) Prediction Error for models $\mathcal{M}_0$ and $\mathcal{M}_4$ on unperturbed evaluation data.

(b) Prediction Error for models $\mathcal{M}_0$ and $\mathcal{M}_1$ on perturbed evaluation data.

Figure 8: Open loop evaluation results.

6.3.4 Real World Driving

We have also evaluated this model on our self-driving car by replacing the existing planner module with the learned model $\mathcal{M}_4$ and have replicated the driving behaviors observed in simulation. The videos of several of these runs are available here and they illustrate not only the smoothness of the network's driving ability, but also its ability to deal with stop-signs and turns and to drive for long durations in full closed-loop control without deviating from the trajectory.

6.4 Open Loop Evaluation

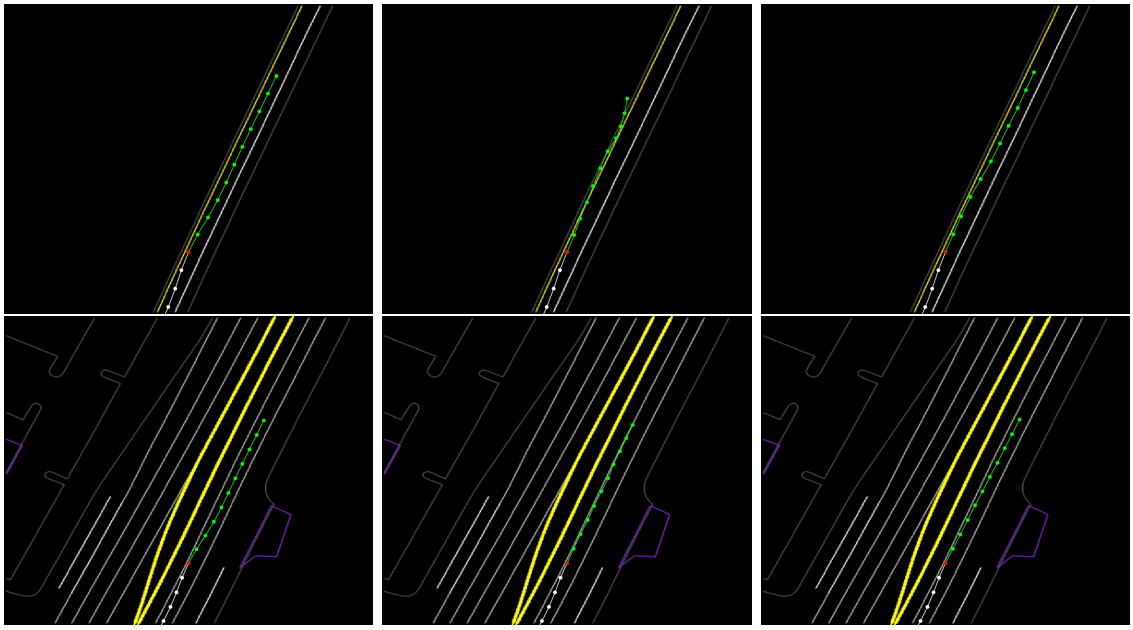
In an open-loop evaluation, we take test examples of expert driving data and for each example, compute the L_2 distance error between the predicted and ground-truth waypoints. Unlike the closed-loop setting, the predictions are not used to drive the agent forward and thus the network never sees its own predictions as input. Fig. 8a shows the L_2 distance metric in this open-loop evaluation setting for models $\mathcal{M}_0$ and $\mathcal{M}_4$ on a test set of 10,000 examples. These results show that model $\mathcal{M}_0$ makes fewer errors than the full model $\mathcal{M}_4$, but we know from closed-loop testing that $\mathcal{M}_4$ is a far better driver than $\mathcal{M}_0$. This shows how open-loop evaluations can be misleading, and closed-loop evaluations are critical while assessing the real performance of such driving models.

We also compare the performance of models $\mathcal{M}_0$ and $\mathcal{M}_1$ on our perturbed evaluation data w.r.t the L_2 distance metric in Fig. 8b. Note that the model trained without including perturbed data ($\mathcal{M}_0$) has larger errors due to its inability to bring the agent back from the perturbation onto its original trajectory. Fig. 9 shows examples of the trajectories predicted by these models on a few representative examples showcasing that the perturbed data is critical to avoiding the veering-off tendency of the model trained without such data.

6.5 Failure Modes

At our ground resolution of 20 cm/pixel, the agent currently sees 64 m in front and 40 m on the sides and this limits the model's ability to perform merges on T-junctions and

Perceptive Exhibit 2047
IPR2025-01576



(a) Ground-truth (b) Model $\mathcal{M}_0$ Prediction (c) Model $\mathcal{M}_1$ Prediction

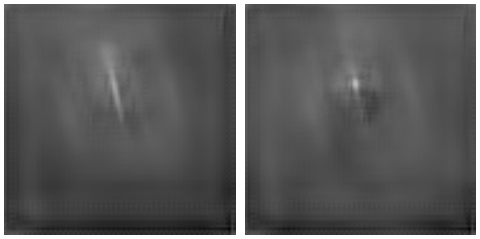
Figure 9: Comparison of ground-truth trajectory in (a) with the predicted trajectories from models $\mathcal{M}_0$ and $\mathcal{M}_1$ in (b) and (c) respectively on two perturbed examples. The red point is the reference pose (u_0, v_0), white points are the past poses and green points are the future poses.

turns from a high-speed road. Specific situations like U-turns and cul-de-sacs are also not currently handled, and will require sampling enough training data. The model occasionally gets stuck in some low speed nudging situations. It sometimes outputs turn geometries that make the specific turn infeasible (e.g. large turning radius). We also see some cases where the model gets over aggressive in novel and rare situations for example by trying to pass a slow moving vehicle. We believe that adequate simulated exploration may be needed for highly interactive or rare situations.

6.6 Sampling Speed Profiles

The waypoint prediction from the model at timestep k is represented by the probability distribution $P_k(x, y)$ over the spatial domain in the top-down coordinate system. In this paper, we pick the mode of this distribution $\mathbf{p}_k$ to update the memory of the $AgentRNN$. More generally, we can also sample from this distribution to allow us to predict trajectories with different speed profiles. Fig. 10 illustrates the predictions $P_1(x, y)$ and $P_5(x, y)$ at the first and the fifth iterations respectively, for a training example where the past motion history has been dropped out. Correspondingly, $P_1(x, y)$ has a high uncertainity along the longitudinal position and allows us to pick from a range of speed samples. Once we pick a specific sample, the ensuing waypoints get constrained in their ability to pick different speeds and this shows as a centered distribution at the $P_5(x, y)$.

Perceptive Exhibit 2047
IPR2025-01576



(a) $\log P_1(x, y)$ (b) $\log P_5(x, y)$

Figure 10: Sampling speed profiles. The probability distribution $P_1(x, y)$ predicted by the model at timestep $k = 1$ allows us to sample different speed profiles conditioned on which the later distribution $P_5(x, y)$ gets more constrained.

The use of a probability distribution over the next waypoint also presents the interesting possibility of constraining the model predictions at inference time to respect hard constraints. For example, such constrained sampling may provide a way to ensure that any trajectories we generate strictly obey legal restrictions such as speed limits. One could also constrain sampling of trajectories to a designated region, such as a region around a given reference trajectory.

7. Discussion

In this paper, we presented our experience with what it took to get imitation learning to perform well in real-world driving. We found that key to its success is synthesizing interesting situations around the expert's behavior and augmenting appropriate losses that discourage undesirable behavior. This constrained exploration is what allowed us to avoid collisions and off-road driving even though such examples were not explicitly present in the expert's demonstrations. To support it, and to best leverage the expert data, we used middle-level input and output representations which allow easy mixing of real and simulated data and alleviate the burdens of learning perception and control. With these ingredients, we got a model good enough to drive a real car. That said, the model is not yet fully competitive with motion planning approaches but we feel that this is a good step forward for machine learned driving models. There is room for improvement: comparing to end-to-end approaches, and investigating alternatives to imitation dropout are among them. But most importantly, we believe that augmenting the expert demonstrations with a thorough exploration of rare and difficult scenarios in simulation, perhaps within a reinforcement learning framework, will be the key to improving the performance of these models especially for highly interactive scenarios.

Acknowledgments

We would like to thank Andrew Barton-Sweeney for help running the model on the car, Aleksandar Gabrovski for help with the simulation, Anca Dragan and Dragomir Anguelov for reviewing the paper and suggesting several improvements.

Perceptive Exhibit 2047
IPR2025-01576

References

Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.

Mariusz Bojarski, Philip Yeres, Anna Choromanska, Krzysztof Choromanski, Bernhard Firner, Lawrence Jackel, and Urs Muller. Explaining how a deep neural network trained with end-to-end learning steers a car. *arXiv preprint arXiv:1704.07911*, 2017.

Chenyi Chen, Ari Seff, Alain Kornhauser, and Jianxiong Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2722–2730, 2015.

Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–9. IEEE, 2018.

Nemanja Djuric, Vladan Radosavljevic, Henggang Cui, Thi Nguyen, Fang-Chieh Chou, Tsung-Han Lin, and Jeff Schneider. Motion prediction of traffic actors for autonomous driving using deep convolutional networks. *arXiv preprint arXiv:1808.05819*, 2018.

Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. *arXiv preprint arXiv:1711.03938*, 2017.

Nathaniel Fairfield and Chris Urmson. Traffic light mapping and detection. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 5421–5426. IEEE, 2011.

Simon Hecker, Dengxin Dai, and Luc Van Gool. Learning driving models with a surround-view camera system and a route planner. *arXiv preprint arXiv:1803.10158*, 2018.

Alex Kuefler, Jeremy Morton, Tim Wheeler, and Mykel Kochenderfer. Imitating driver behavior with generative adversarial networks. In *Intelligent Vehicles Symposium (IV), 2017 IEEE*, pages 204–211. IEEE, 2017.

Michael Laskey, Caleb Chuck, Jonathan Lee, Jeffrey Mahler, Sanjay Krishnan, Kevin Jamieson, Anca Dragan, and Ken Goldberg. Comparing human-centric and robot-centric sampling for robot deep learning from demonstrations. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 358–365. IEEE, 2017a.

Michael Laskey, Jonathan Lee, Roy Fox, Anca Dragan, and Ken Goldberg. Dart: Noise injection for robust imitation learning. *arXiv preprint arXiv:1703.09327*, 2017b.

Donghan Lee, Youngwook Paul Kwon, Sara McMains, and J Karl Hedrick. Convolution neural network-based lane change intention prediction of surrounding vehicles for acc. In *Intelligent Transportation Systems (ITSC), 2017 IEEE 20th International Conference on*, pages 1–6. IEEE, 2017.

Perceptive Exhibit 2047
IPR2025-01576

Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2018.

Matthias Müller, Alexey Dosovitskiy, Bernard Ghanem, and Vladen Koltun. Driving policy transfer via modularity and abstraction. *arXiv preprint arXiv:1804.09364*, 2018.

Brian Paden, Michal Čáp, Sze Zheng Yong, Dmitry Yershov, and Emilio Frazzoli. A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on intelligent vehicles*, 1(1):33–55, 2016.

Xinlei Pan, Yurong You, Ziyan Wang, and Cewu Lu. Virtual to real reinforcement learning for autonomous driving. *arXiv preprint arXiv:1704.03952*, 2017.

Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. Deepxplore: Automated whitebox testing of deep learning systems. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 1–18. ACM, 2017.

Dean A Pomerleau. Alvinn: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pages 305–313, 1989.

Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635, 2011.

Axel Sauer, Nikolay Savinov, and Andreas Geiger. Conditional affordance learning for driving in urban environments. *arXiv preprint arXiv:1806.06498*, 2018.

Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.

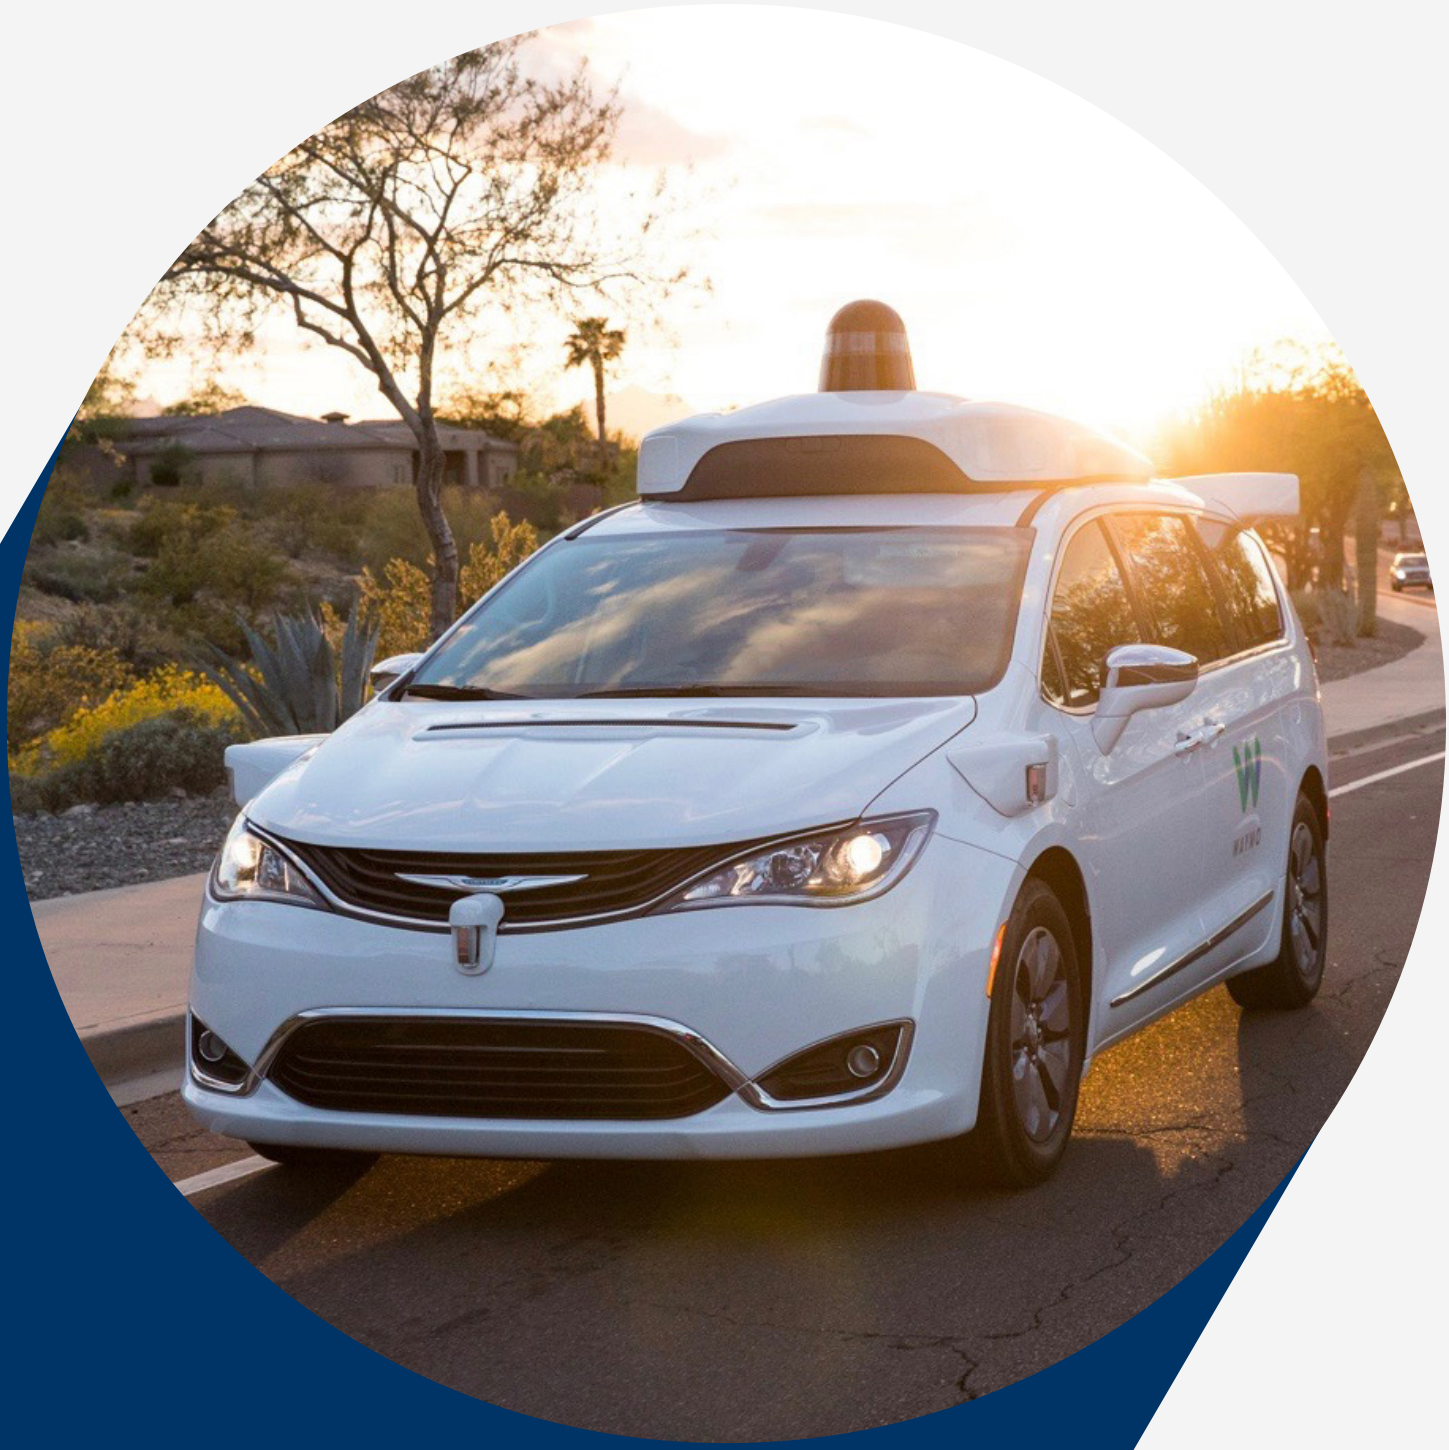
Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. Deeptest: Automated testing of deep-neural-network-driven autonomous cars. In *Proceedings of the 40th International Conference on Software Engineering*, pages 303–314. ACM, 2018.

Huazhe Xu, Yang Gao, Fisher Yu, and Trevor Darrell. End-to-end learning of driving models from large-scale video datasets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2174–2182, 2017.

Bin Yang, Ming Liang, and Raquel Urtasun. Hdnet: Exploiting hd maps for 3d object detection. In *Conference on Robot Learning*, pages 146–155, 2018.

Perceptive Exhibit 2047
IPR2025-01576

Attachment D



ON THE ROAD TO FULLY SELF-DRIVING

Waymo Safety Report

WAYMO

Perceptive Exhibit 2047
IPR2025-01576

OUR MISSION

Waymo's mission is to bring self-driving technology to the world, making it safe and easy for people and things to move around. We believe our technology can improve mobility by giving people the freedom to get around, and save thousands of lives now lost to traffic crashes.

WAYMO



We're Building a Safer Driver for Everyone





Self-driving vehicles hold the promise to improve road safety and offer new mobility options to millions of people. Whether they're saving lives or helping people run errands, commute to work, or drop kids off at school, fully self-driving vehicles hold enormous potential to transform people's lives for the better.

Safety is at the core of Waymo's mission—it's why we were founded over eight years ago as the *Google self-driving car project*.

Every year, 1.2 million lives are lost to traffic crashes around the world, and in the U.S. the number of tragedies is growing. A common element of these crashes is that 94% involve human error. Driving is not as safe or as easy as it should be, while distracted driving is on the rise. We believe our technology could save thousands of lives now lost to traffic crashes every year.

Our commitment to safety is reflected in everything we do, from our company culture to how we design and test our technology. In this, our first Safety Report on Waymo's fully self-driving technology, we detail Waymo's work on—and our commitment to—safety. This overview of our safety program underscores the important lessons learned through the 3.5 million miles Waymo's vehicles have self-driven on public roads and through our billions of miles of simulated driving.

Perceptive Exhibit 2047
IPR2025-01576





Waymo's Safety Report also addresses the U.S. Department of Transportation (DOT) federal policy framework for autonomous vehicles: *Automated Driving Systems 2.0: A Vision for Safety*. The DOT framework outlines 12 safety design elements, and encourages companies testing and deploying self-driving systems to address each of these areas. Over the course of our Report, we will outline the processes relevant to each safety design element and how they underpin the development, testing, and deployment of fully self-driving vehicles.

Fully self-driving vehicles will succeed in their promise and gain public acceptance only if they are safe. That's why Waymo has been investing in safety and building the processes that give us the confidence that our self-driving vehicles can serve the public's need for safer transportation and better mobility.



Perceptive Exhibit 2047
IPR2025-01576

TABLE OF CONTENTS

1. Our System Safety Program: Safety by Design 10

Areas Addressed by Waymo's System Safety Program 11

Safety Processes 12

2. How Waymo's Self-Driving Vehicles Work 13

The Self-Driving System 13

Object and Event Detection and Response: Our Vehicle Sensors 14

Our Self-Driving Software 15

Operational Design Domain: Ensuring Our Vehicles Operate Safely Under Specific Conditions 16

Minimal Risk Condition (Fallback): Ensuring the Vehicle Can Transition to a Safe Stop 17

Data Recording and Post-Crash Behavior 18

Self-Driving Vehicle Cybersecurity 18

3. Testing and Validation Methods: Ensuring Our Vehicles Are Capable and Safe 20

Base Vehicle Safety 21

Self-Driving Hardware Testing 21

Self-Driving Software Testing 22

Behavioral Competencies for Normal Driving 25

Testing the Fully Integrated Self-Driving Vehicle 26

Testing on Public Roads 26

Testing Crash Avoidance Capabilities 28

Hardware Reliability and Durability Testing 29

4. Interacting Safely with the Public 30

Rider Experience 31

Accessibility: Unlocking Opportunities for Those Who Cannot Drive Today 32

Emergencies and Interacting With Law Enforcement and First Responders 33

5. Conclusion 34

Appendix: Scenario Types Used for Testing and Validation 35

Appendix A. Basic Behavioral Competency Testing 36

Appendix B. Avoidance or Mitigation of Common Crash Scenarios 38

Glossary 40

End Notes 42

Perceptive Exhibit 2047
IPR2025-01576

The World Around Us

SAFETY

94
PERCENT
of U.S. crashes involve human error. [1]

1.2
MILLION
deaths worldwide due to vehicle crashes in 2013. [2]

37,461
ROAD DEATHS
in the U.S. in 2016 and 2.4 million injuries in 2015. [3]

2
OUT OF 3
people will be involved in a drunk driving crash in their lifetime. [4]

SOCIETY

$594
BILLION
in harm from loss of life and injury each year. [5]

$277
BILLION
in annual economic costs. [6]

$160
BILLION
in gas burned and time lost each year. [7]

MOBILITY AND QUALITY OF LIFE

3
MILLION
Americans age 40 and older are blind or have low vision. [8]

79
PERCENT
of seniors age 65 and older living in car-dependent communities. [9]

42
HOURS
wasted in traffic each year per person. [7]

Self-driving Technology Can Save Lives and Improve Mobility

WAYMO

At Waymo, we're designing fully self-driving vehicles that make it safe and easy for everyone to get around.

Perceptive Exhibit 2047
IPR2025-01576

Building the World's Most Experienced Driver

Every mile, in every car, is shared with the entire fleet, giving every Waymo vehicle more experience for the next mile.



4 Generations of Self-Driving Vehicles

8 Years Self-Driving in More than 20 U.S. Cities

3.5 Million Real-World Miles on Public Roads

2.5 Billion Self-Driven Miles Simulated in 2016

Perceptive Exhibit 2047
IPR2025-01576

How Our Self-Driving Vehicle Sees the World and How it Works

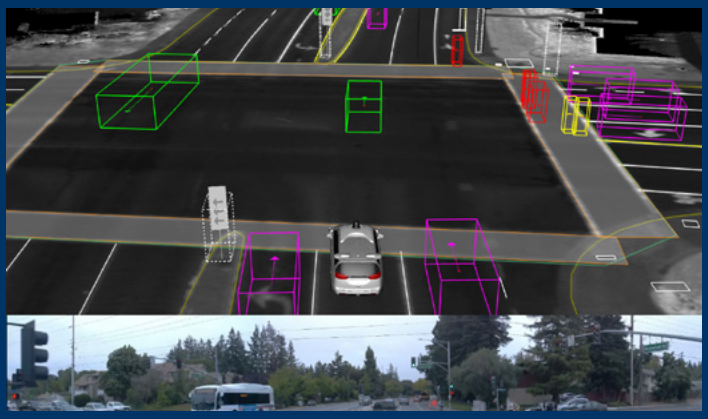
At the most basic level, human drivers need to answer four questions: "Where am I?" (perceiving the environment around you), "What's around me?" (processing that information), "What will happen next?" (predicting how others in that environment will behave), and "What should I do?" (making driving decisions based on that information). Self-driving vehicles need to answer those questions, too.



A Waymo vehicle's onboard map view of the intersection of W. Middlefield Road at Rengstorff Avenue, Mountain View, CA.

1. Where Am I?

Before our cars drive in any location, our team builds our own detailed three-dimensional maps that highlight information such as road profiles, curbs and sidewalks, lane markers, crosswalks, traffic lights, stop signs, and other road features. Rather than rely on GPS, Waymo's vehicles cross-reference their pre-built maps with real-time sensor data to precisely determine their location on the road.



In this example, our vehicle has detected vehicles (depicted by green and purple boxes), pedestrians (in yellow), and cyclists (in red) at the intersection—and a construction zone up ahead.

2. What's Around Me?

Our sensors and software scan constantly for objects around the vehicle—pedestrians, cyclists, vehicles, road work, obstructions—and continuously read traffic controls, from traffic light color and railroad crossing gates to temporary stop signs. Our vehicles can see up to 300 meters away (nearly three football fields) in every direction.

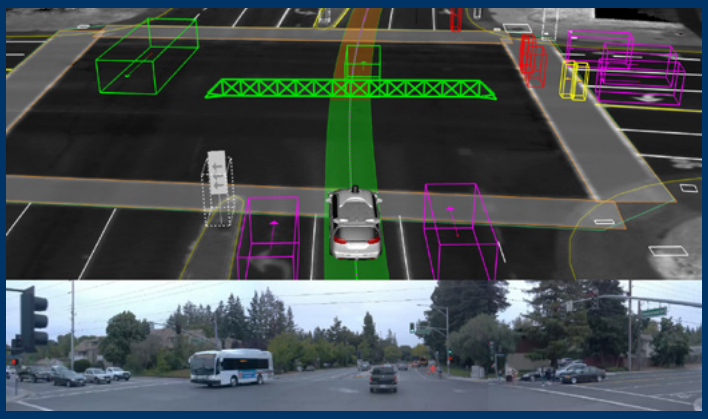
Perceptive Exhibit 2047
IPR2025-01576



The simulated imagery shown demonstrates how our software assigns predictions to each object surrounding our vehicle—other vehicles, cyclists, pedestrians, and more.

3. What Will Happen Next?

For every dynamic object on the road, our software predicts future movements based on current speed and trajectory. It understands that a vehicle will move differently than a cyclist or pedestrian. The software then uses that information to predict the many possible paths that other road users may take. Our software also takes into account how changing road conditions (such as a blocked lane up ahead) may impact the behavior of others around it.



The green path indicates the trajectory through which our vehicle can proceed ahead. The series of green fences indicate that the self-driving vehicle can proceed, and that the vehicle has identified the vehicles ahead and understands it has to maintain certain headway.

4. What Should I Do?

The software considers all of this information as it finds an appropriate route for the vehicle to take. Our software selects the exact trajectory, speed, lane, and steering maneuvers needed to progress along this route safely. Because our vehicles are constantly monitoring the environment, and predicting the future behavior of other road users in 360 degrees around our vehicles, they're able to respond quickly and safely to any changes on the road.

Perceptive Exhibit 2047
IPR2025-01576



SECTION 1

Our System Safety Program

SAFETY BY DESIGN

As the first company to complete a fully self-driving trip on public roads, Waymo has had to write its own playbook.

In the earliest days of our company, we established our System Safety Program, which documented practices that would ensure safety in the testing and development of our technology. Today, that program is a comprehensive and robust approach we call *Safety by Design*.

Safety by Design means we consider safety from the ground up and incorporate safety at every system level and every development stage, from design to testing and validation. It is a multi-pronged approach that builds upon best practices from a variety of industries, including aerospace, automotive, and defense (including aspects of MIL-STD-882E [10] and ISO 26262). [11]

In line with these practices, each individual component of our self-driving vehicle is tested robustly to ensure that all subsystems perform safely when integrated as a complete self-driving system. This approach also helps us validate that the vehicle works safely as a fully self-driving vehicle on the road, and understand how a change or failure in any part of the system—component, subsystem, or otherwise—causes changes throughout the rest of the self-driving system.

This process has led to many of Waymo's key safety features, including redundant critical safety systems, which enable the vehicle to come to a safe stop in the event of a technology failure, the use of complementary sensors with overlapping fields-of-view, and our extensive testing program which has helped us make rapid improvements in our technology.

Perceptive Exhibit 2047
IPR2025-01576

Areas Addressed by Waymo's System Safety Program

Our System Safety Program addresses five distinct safety areas: behavioral safety, functional safety, crash safety, operational safety, and non-collision safety. Each aspect requires a combination of testing methods that, taken together, allow us to validate the safety of our fully self-driving vehicles.

Behavioral Safety:
Behavioral safety refers to the driving decisions and behavior of our vehicles on the road. Just as for human drivers, our vehicles are subject to traffic rules and must safely navigate a variety of scenarios, both expected and unexpected. Waymo uses a combination of functional analysis, simulation tools, and on-road driving to fully understand the challenges presented within our operational design domain, and to develop safety requirements and a multi-pronged testing and validation process.

Functional Safety:
Functional safety seeks to ensure that our vehicles operate safely even when there is a system fault or failure. That means building in backup systems and redundancies to handle the unexpected. For example, all of our self-driving vehicles are equipped with a secondary computer that can take over in the event of a main computer failure, bringing the vehicle to a safe stop (i.e. a minimal risk condition). Each of our vehicles also has backup steering and braking, along with many layers of redundancies throughout the system.

Crash Safety:
Crash safety, or crashworthiness, refers to the ability of vehicles to protect passengers inside the vehicles through a variety of measures, ranging from a structural design that shields people inside, to features like seat restraints and airbags that mitigate injury or prevent death. Crash safety in the U.S. is covered by the Federal Motor Vehicle Safety Standards (FMVSS), which are issued by the National Highway Traffic Safety Administration (NHTSA). Vehicle manufacturers must certify that their base vehicles meet applicable FMVSS requirements.

Operational Safety:
This refers to the interaction between our vehicles and passengers. With operational safety, we can ensure that consumers have a safe and comfortable experience in our vehicles. Our approach to building a safe product is informed by our hazard analyses, existing safety standards, extensive testing, and best practices from a variety of industries. For example, through initiatives like our early rider program (further described in section 4), we have developed and tested user interfaces so that passengers can clearly indicate their destination, direct the vehicle to pull over, and contact Waymo rider support.

Non-Collision Safety:
We address physical safety for the range of people who might interact with the vehicle. For example, this includes electrical system or sensor hazards that could cause harm to occupants, vehicle technicians, test drivers, first responders, or bystanders.

Perceptive Exhibit 2047
IPR2025-01576

Safety Processes

Waymo organizes the processes we use to keep our vehicles safe through our System Safety Program. Safety requirements needed to reduce the risk of potential hazards are captured internally, addressed in design, and then verified and validated to demonstrate that safety risks have been reduced to the levels identified in the analyses.

Our approach starts with identifying hazard scenarios and potential mitigations that can be implemented to reduce risk. These mitigations may take various forms such as software or hardware requirements, hardware or software design recommendations, procedural controls, or recommendations for additional analyses. We use various hazard assessment methods such as preliminary hazard analysis, fault tree, and Design Failure Modes and Effects Analyses (DFMEA). This continuous process goes hand-in-hand with ongoing engineering and test activities and safety engineering analyses.

The hazard analysis process helps identify requirements for our self-driving system architecture, subsystems, and components. These safety requirements are developed from the use of a series of subsystem and system analysis techniques, various systems engineering processes, and Federal and State laws and regulations. The analysis also supports the development of requirements for our behavioral safety testing, and how our system detects and handles faults.

With our system architecture and requirements defined, Waymo then conducts extensive testing on public roads, a closed course, and in simulated driving. We use information gathered from this testing, as well as research into national crash data and naturalistic driving studies [12], to provide additional insights into potential hazards. The combined knowledge derived from these various tools plays a major role in our understanding of our system's readiness. Drawing on this understanding, we're able to comprehensively analyze and evaluate the safety of our system before we permit fully self-driving operation on public roads.



Perceptive Exhibit 2047
IPR2025-01576

SECTION 2



How Waymo's Self-Driving Vehicles Work

The Case for Full Autonomy: Allowing Passengers to Stay Passengers

Advanced driver-assist technologies were one of the first technologies our teams explored. In 2012 we developed and tested a Level 3 system that would drive autonomously on the freeway in a single lane but would still require a driver to take over at a moment's notice. During our internal testing, however, we found that human drivers over-trusted the technology and were not monitoring the roadway carefully enough to be able to safely take control when needed.

As driver-assist features become more advanced, drivers are often asked to transition from passenger to driver in a matter of seconds, often in challenging or complex situations with little context of the scene ahead. The more tasks the vehicle is responsible for, the more complicated and vulnerable this moment of transition becomes.

Avoiding this "handoff problem" is part of the reason why Waymo is working on *fully* self-driving vehicles. Our technology takes care of all of the driving, allowing passengers to stay passengers.



The Self-Driving System

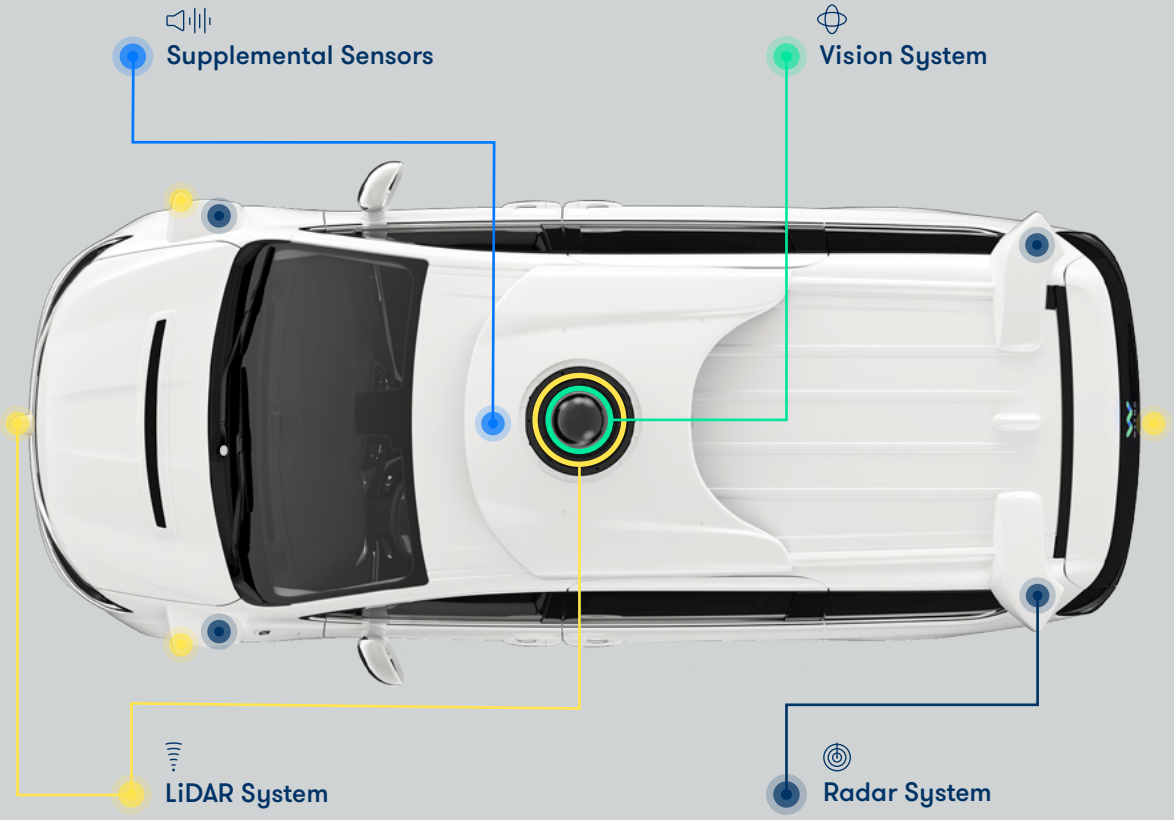
Our fully self-driving system is designed to operate without a human driver, unlike technologies sold in cars today such as adaptive cruise-control or lane-keeping systems which require constant monitoring by the driver. Our system includes the software and hardware that, when integrated into the vehicle, perform all driving functions.

In self-driving jargon, Waymo's self-driving system is designed to perform the entire dynamic driving task within a geographic area and under certain defined conditions, without the need for a human driver. This type of technology falls under SAE International's definition of a Level 4 automated driving system, as our technology also has the ability to bring a vehicle to a safe stop (i.e. a minimal risk condition) in the event of any system failures. Unlike autonomous systems at lower levels (Level, 1, Level 2, and Level 3), a Level 4 system also has the ability to bring a vehicle to a safe stop (i.e. achieve a minimal risk condition) in the event of any system failures, without any expectation that a human driver take over. [13]

Perceptive Exhibit 2047
IPR2025-01576

Object and Event Detection and Response: Our Vehicle Sensors

To meet the complex demands of autonomous driving, Waymo has developed an array of sensors that allow our vehicle to see 360 degrees, both in daytime and at night, and up to nearly three football fields away. This multi-layered sensor suite works together seamlessly to paint a detailed 3D picture of the world, showing dynamic and static objects including pedestrians, cyclists, other vehicles, traffic lights, construction cones, and other road features.



LiDAR (Laser) System

LiDAR (Light Detection and Ranging) works day and night by beaming out millions of laser pulses per second—in 360 degrees—and measuring how long it takes to reflect off a surface and return to the vehicle. Waymo's system includes three types of LiDAR developed in-house: a short-range LiDAR that gives our vehicle an uninterrupted view directly around it, a high-resolution mid-range LiDAR, and a powerful new generation long-range LiDAR that can see almost three football fields away.

Vision (Camera) System

Our vision system includes cameras designed to see the world in context, as a human would, but with a simultaneous 360-degree field of view, rather than the 120-degree view of human drivers. Because our high-resolution vision system detects color, it can help our system spot traffic lights, construction zones, school buses, and the flashing lights of emergency vehicles. Waymo's vision system is comprised of several sets of high-resolution cameras, designed to work well at long range, in daylight and low-light conditions.

Radar System

Radar uses wavelengths to perceive objects and movement. These wavelengths are able to travel around objects like rain drops, making radar effective in rain, fog, and snow, day or night. Waymo's radar system has a continuous 360-degree view, so it can track the speed of road users in front, behind and to both sides of the vehicle.

Supplemental Sensors

Waymo vehicles also have a number of additional sensors, including our audio detection system that can hear police and emergency vehicle sirens up to hundreds of feet away, and GPS to supplement our vehicles' extensive understanding of their physical locations in the world.

Perceptive Exhibit 2047
IPR2025-01576

Our Self-Driving Software

Our self-driving software is the "brain" of our vehicle. It makes sense of the information coming from our sensors, and uses that information to make the best driving decisions for each situation.

Waymo has spent eight years building and refining our software, using machine learning and other advanced engineering techniques. We've trained our software through years of careful design and testing, billions of miles of simulated driving, and more than 3.5 million miles of on-road driving experience.

Our system possesses a deep, contextual understanding of the world; this is a key part of what differentiates Level 4 technology. Our self-driving software doesn't just detect the presence of other objects; it actually understands what an object is, how it's likely to behave, and how that should affect our vehicle's own behavior on the road. This is how our vehicles safely navigate roads in fully autonomous mode.

While our software is made up of many different pieces, here we detail three main components: perception, behavior prediction, and planner.

PERCEPTION

Perception is the part of our software that detects and classifies objects on the road, while also estimating their speed, heading, and acceleration over time. Our self-driving software takes the myriad of details coming from Waymo's sensors and turns them into a cohesive real-time view of the world. Perception helps our vehicle distinguish pedestrians, cyclists, motorcyclists, vehicles, and more. It also distinguishes the color of static objects such as traffic signals. For these kinds of objects, perception enables our system to semantically understand the situation around our vehicle—whether a light is green and clear for the vehicle to proceed, or whether a lane is blocked because of the many cones in front of it.

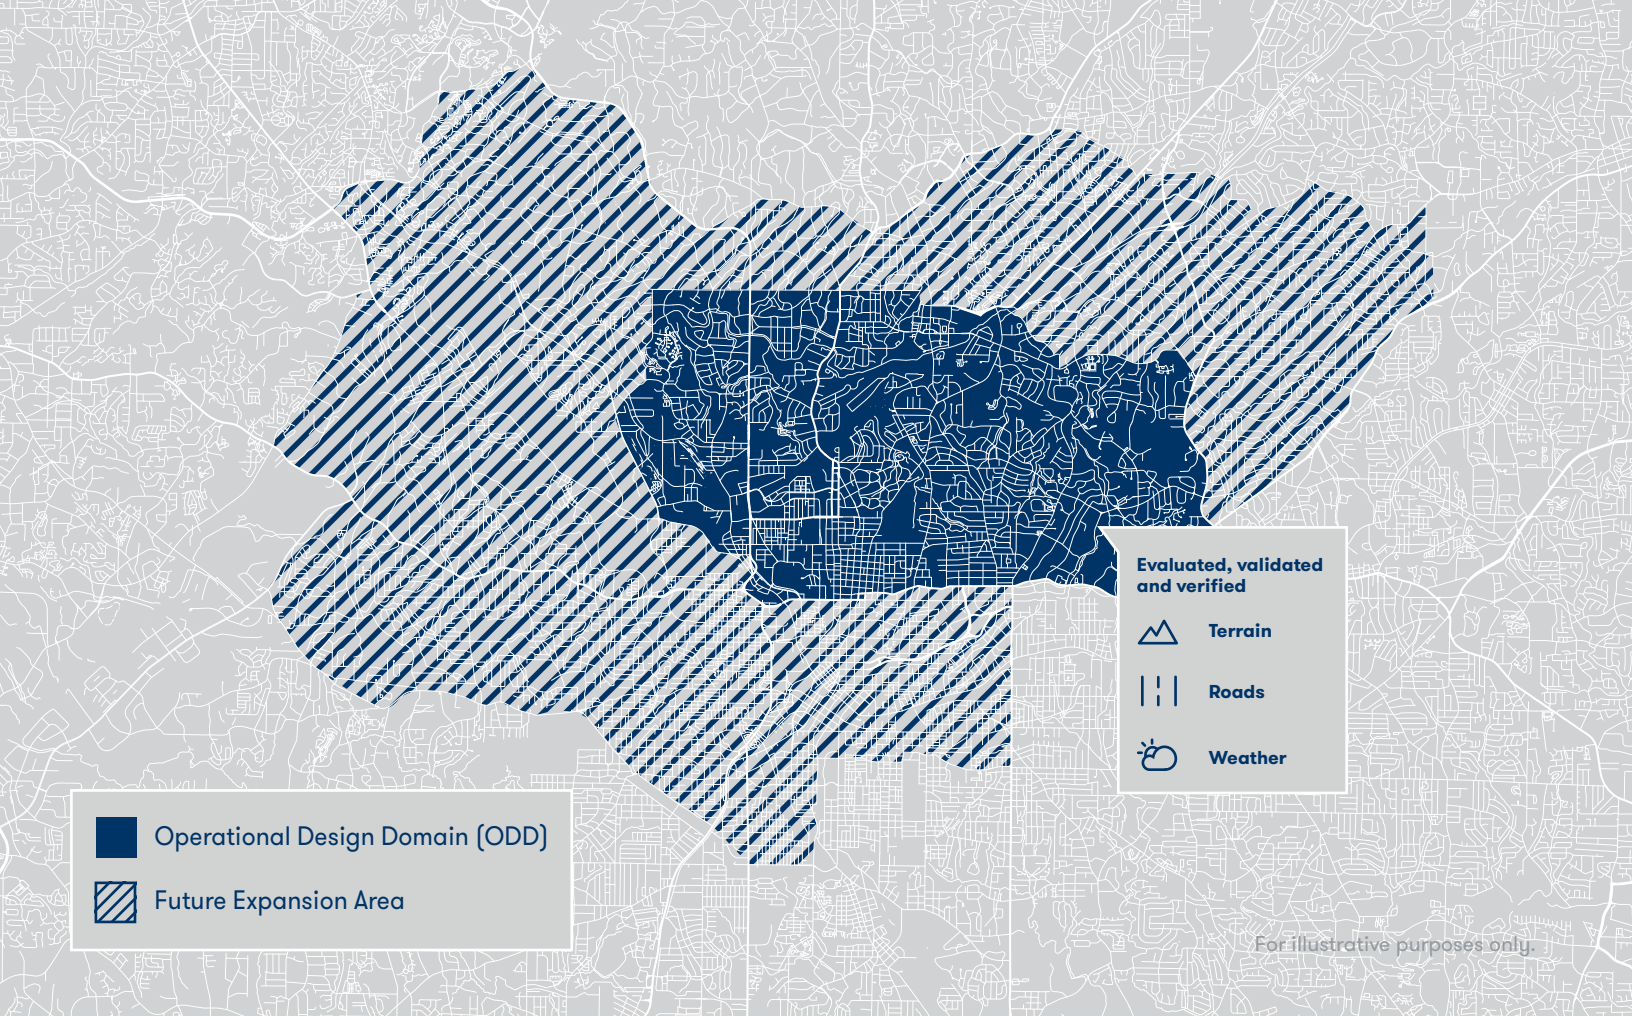
BEHAVIOR PREDICTION

With *behavior prediction*, our software can model, predict, and understand the intent of each object on the road. Because Waymo has millions of miles of driving experience, our vehicles have highly accurate models of how different road users are likely to behave. For example, our software understands that, though pedestrians, cyclists, and motorcyclists may look similar, their behavior can vary dramatically. Pedestrians move more slowly than either cyclists or motorcyclists, but they can change direction more suddenly.

PLANNER

Our *planner* considers all the information our software has gathered from perception and behavior prediction, and plots out a path for our vehicles. In our experience, the best drivers are the defensive drivers. That's why we've baked in defensive driving behaviors, such as staying out of other drivers' blind spots and leaving extra room for cyclists and pedestrians. Waymo's planner can also think several steps ahead. For example, if our software perceives that an adjacent lane ahead is closed due to construction, and predicts that a cyclist in that lane will move over, our planner can make the decision to slow down or make room for the cyclist well ahead of time. Using our on-road experience, we're also refining our driving so our movements on the road are smooth and comfortable for passengers inside our vehicles, and natural and predictable for other road users.

Perceptive Exhibit 2047
IPR2025-01576



***Operational Design Domain:* Ensuring Our Vehicles Operate Safely Under Specific Conditions**

The operational design domain refers to the conditions under which a self-driving system can safely operate. Waymo's domain includes geographies, roadway types, speed range, weather, time of day, and state and local traffic laws and regulations.

An operational design domain can be very limited: for instance, a single fixed route on low-speed public streets or private grounds (such as business parks) in temperate weather conditions during daylight hours. However, Waymo aims to have a broad operational design domain to cover everyday driving. We're developing self-driving technology that can navigate city streets in a variety of conditions within broad geographic areas. Our vehicles are designed with the capability to drive in inclement weather, such as light to moderate rain, and can operate in daytime and at night.

Waymo's system is also designed so each vehicle does not operate outside of its approved operational design domain. For example, passengers cannot select a destination outside of our approved geography, and our software will not create a route that travels outside of a "geo-fenced" area, which has been mapped in detail (see "How We Build a Map for a Self-Driving Vehicle"). Similarly, our vehicles are designed to automatically detect sudden changes (such as a snowstorm) that would affect safe driving within their operational design domain and come to a safe stop (i.e. achieve a "minimal risk condition") until conditions improve.

We design our vehicles to be capable of complying with federal, state, and local laws within their geographic area of operations. [14] Legal requirements, and any changes in those requirements, are identified and built into our system as safety requirements, including relevant speed limits, traffic signs, and signals. Before our vehicles drive in a new location, our team works to understand any unique road rules or driving customs, and we update our software so our vehicles are capable of responding safely. For example, California and Texas (home to two of Waymo's test cities) have differing rules for how to make right turns in the presence of a bike lane.

Waymo's operational design domain continues to evolve. Our ultimate goal is to develop fully self-driving technology that can take someone from A to B, anytime, anywhere, and in all conditions. As our system's capabilities grow and are validated, we will expand our operational design domain to bring our technology to more people.

Perceptive Exhibit 2047
IPR2025-01576

Minimal Risk Condition (Fallback): Ensuring the Vehicle Can Transition to a Safe Stop
Vehicles with lower levels of automation rely on a human driver to take back control if a situation on the road becomes too complex for the technology to handle, or if the technology itself fails. As a fully self-driving system, Waymo's technology must be robust enough to handle these situations on its own.

If our self-driving vehicle can no longer proceed on a planned trip, it must be capable of performing a safe stop, known as a "minimal risk condition" or fallback. This might include situations when the self-driving system experiences a problem, when the vehicle is involved in a collision, or when environmental conditions change in a way that would affect safe driving within our operational design domain.

Waymo's system is designed to detect each one of these scenarios automatically. In addition, our vehicles run thousands of checks on their systems every second, looking for faults. Our system is equipped with a series of redundancies for critical systems, such as sensors, computing, and braking. How our vehicle responds varies with the type of roadway on which a situation occurs, the current traffic conditions, and the extent of the technology failure. Depending on these factors, the system will determine an appropriate response to keep the vehicle and its passengers safe, including pulling over or coming to a safe stop. [15]

Our Vehicles' Redundant Safety-Critical Systems



Backup Computing

A secondary computer in the vehicle is always running in the background and is designed to bring the vehicle to a safe stop if it detects a failure of the primary system.



Backup Braking

If the primary braking system fails, we have a full secondary braking system that immediately kicks in. Either braking system can bring the vehicle to a safe stop if a failure occurs in the other.



Backup Steering

The steering system features a redundant drive motor system with independent controllers and separate power supplies. Either one can manage steering in the case that a failure occurs in the other.



Backup Power Systems

Independent power sources are provided for each of the critical driving systems. These independent power sources ensure that our vehicles' critical driving components remain online during single power failures or circuit interruptions.



Backup Collision Detection and Avoidance System

Multiple backup systems—including independent collision avoidance systems—constantly scan the road immediately ahead and behind the vehicle for objects such as pedestrians, cyclists, and other vehicles. These redundant systems slow or stop the vehicle in the rare event that the primary system does not detect or respond to objects in the path of the vehicle.



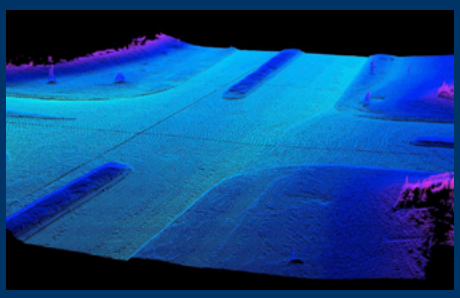
Redundant Inertial Measurement Systems for Vehicle Positioning

Redundant inertial measurement systems help the vehicle accurately track its motion along the road. These two systems cross-check each other and assume control from one another, if a fault is detected in either system.

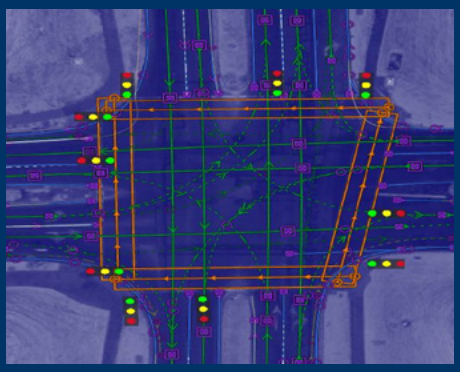
Perceptive Exhibit 2047
IPR2025-01576

How We Build a Map for a Self-Driving Vehicle

Before our vehicles are introduced on the road, our mapping team first uses our sensors on test vehicles to create highly detailed 3D maps. These maps are distinct from basic satellite imagery or online maps. Instead, Waymo's maps provide our vehicle with a deep understanding of the physical environment: road types, the distance and dimensions of the road itself, and other topographical features.



We take this data and add salient information that includes traffic control information such as the lengths of crosswalks, the locations of traffic lights, and relevant signage.



With our maps installed onboard our vehicle, our system can then focus on the parts of the environment that change dynamically around it, such as other road users. Our system can detect when a road has changed by cross-referencing the real-time sensor data with its on-board 3D map. If a change in the roadway (e.g., a collision up ahead that closes an intersection) is detected, our vehicle can re-route itself within the system's operational design domain and alert our operations center so that other vehicles in the fleet can avoid the area. In this case, the maps not only serve as an added reference point to our software, but also provide important feedback to our system.

These detailed custom maps give us a comprehensive understanding of the conditions in every location where we drive. When coupled with our deep knowledge of the capabilities of our system, they help us ensure that our vehicles operate only within their operational design domain.

Data Recording and Post-Crash Behavior

Waymo's self-driving technology never stops improving. Waymo has a robust system for collecting and analyzing data from encounters we have on the road. [16] Anything we learn from the experience of one vehicle, we apply to our entire fleet.

Waymo's system can detect when it has been involved in a collision and will notify our Waymo operations center automatically. There our trained specialists can initiate post-crash procedures, which include procedures for interacting with law enforcement and first responders, and for sending members of our team on location. Our operations center also has rider support specialists, who can communicate directly with our passengers through our in-vehicle audio system.

Following a collision, we're able to analyze all available data, including video and other sensor data, to evaluate factors that may have contributed to the incident, and we're able to make any appropriate software changes and update every vehicle in our fleet accordingly. Any damage our vehicles sustain in a collision is repaired and the vehicles are tested for safety before they return to the road.

Self-Driving Vehicle Cybersecurity

Waymo has developed a robust process to identify, prioritize, and mitigate cybersecurity threats. Our security practices are built on the foundation of Google's Security processes and are informed by publications like the NHTSA Cybersecurity Guidance and the Automotive Information Sharing and Analysis Center's (Auto-ISAC) Automotive Cybersecurity Best Practices. To help develop future security best practices, Waymo has also joined the Auto-ISAC, an industry-operated initiative created to enhance cybersecurity awareness and collaboration across the global automotive industry.

Perceptive Exhibit 2047
IPR2025-01576

Waymo's Approach to Security

1. Build Verifiable Software and Systems
2. Encrypt and Verify Channels of Communication
3. Build Redundant Security Measures for Critical Systems
4. Limit Communication Between Critical Systems
5. Provide Timely Software Updates
6. Model and Prioritize Threats

We complete a comprehensive review of all potential security access points to our self-driving system from both the interior and exterior of the physical vehicle, and take steps to limit the number and function of those access points.

This begins by collaborating with our OEM partners at the onset to identify and mitigate vulnerabilities of the base vehicle. Our software and vehicle design processes take full account of known threats to ensure that our system and vehicle designs are protected against them. New software releases go through an extensive peer review and verification process. Our hazard analysis and risk assessment processes have been designed to identify and mitigate risks that might affect safety, including those related to cybersecurity.

In our design, safety-critical aspects of Waymo's vehicles—e.g. steering, braking, controllers—are isolated from outside communication. For example, both the safety-critical computing that determines vehicle movements and the onboard 3D maps are shielded from, and inaccessible from, the vehicle's wireless connections and systems.

We also consider the security of our wireless communication. Our vehicles do not rely on a constant connection to operate safely. While on the road, all communications (e.g., redundant cellular connections) between the vehicles and Waymo are encrypted, including those between Waymo's operations support staff and our riders. Our vehicles can communicate with our operations center to gather more information about road conditions, while our vehicles maintain responsibility for the driving task at all times.

These protections help prevent anyone with limited physical access to our self-driving vehicles, whether passengers or malicious actors nearby, from impairing or altering their security. We have diverse mechanisms for noticing anomalous behavior and internal processes for analyzing those occurrences. Should we become aware of an indication that someone has attempted to impair our vehicle's security, Waymo will trigger its company-wide incident response procedure, which involves impact assessment, containment, recovery, and remediation.

Perceptive Exhibit 2047
IPR2025-01576

SECTION 3

Testing and Validation Methods

ENSURING OUR VEHICLES ARE CAPABLE AND SAFE



Waymo's technology undergoes extensive testing—on the road, in closed courses, and in simulation—so that every part of our system is capable, reliable, and safe when operating within its design domain.

Waymo's self-driving vehicles consist of three primary subsystems that are individually and rigorously tested:

1. The base vehicle, as certified by the OEM
2. Our in-house hardware, including sensors and computers
3. Our self-driving software that makes all the driving decisions

Each of these subsystems is then combined to form a fully integrated self-driving vehicle, which is then further tested and validated. Collectively testing the hardware and software ensures that our overall self-driving vehicle meets all the safety requirements that we have set for our system.

Perceptive Exhibit 2047
IPR2025-01576

Building a Safe and Reliable Supply Chain

Waymo works directly with our suppliers and subcontractors on the performance, safety, quality, and reliability requirements of system components. We include those suppliers in the Failures Modes and Effects Analysis (FMEA) and risk assessment processes as we work to identify potential risks associated with manufacturing processes, stand-alone components, or components when integrated with other subsystems. We monitor the performance of components in our products during manufacturing, and conduct ongoing reliability testing to make sure they meet design expectations and safety requirements before being integrated into our vehicles.

Testing At Every Level

Our multi-layered approach to safety is influenced by the same systems engineering processes used by NASA to launch the *Mars Rover* – a self-driving vehicle that operates millions of miles from Earth.

This approach means we analyze and test our system at the lowest component level to ensure the performance and reliability of our most critical systems. For example, our vision (camera) system alone is subjected to over a hundred separate tests in our labs before a single vehicle with this technology leaves our garage.

First we perform tests on every component: we examine the individual cameras that will make up the vision system, as well as parts like cables and connectors, to determine that each works to designed specifications. Then we test again once the cameras are fully integrated. Each camera is assembled into a ring formation and we calibrate them to work together, checking that the angle and orientation of each camera combines to give a complete 360-degree view. Then we test the vision system as a whole. The assembled camera ring is added to the self-driving system and our engineers run tests to ensure our different sensors work cohesively as one.

Before we operate a single vehicle on public roads with this new vision system, we perform another level of tests, confirming that the vision system is doing its job, performing tasks like seeing traffic lights in a variety of lighting conditions, detecting pedestrians, and spotting construction cones.

Only then is this vision system, as part of our self-driving vehicle, ready for the road.

Base Vehicle Safety

Waymo's current generation self-driving vehicle is a modified version of the 2017 Chrysler Pacifica Hybrid Minivan, into which we have integrated our self-driving system. The modified 2017 Chrysler Pacifica Hybrid Minivans that Fiat Chrysler Automobiles ("FCA") has sold to us have been certified by the manufacturer as compliant with all applicable Federal Motor Vehicle Safety Standards (FMVSS), which standards regulate the safety performance requirements for motor vehicles or items of motor vehicle equipment in the U.S.

Self-Driving Hardware Testing

Through a technical collaboration between FCA and Waymo, we engineered and integrated Waymo's self-driving system, including our self-driving sensors and hardware, with the modified Chrysler Pacifica Hybrid Minivans provided by FCA. To ensure that we have properly integrated our self-driving system into the Chrysler Pacifica Hybrid Minivans that make up our fleet, Waymo has performed thousands of additional tests on top of those completed by FCA. These tests are completed at our private test tracks, in our labs, and in simulation, and are used to evaluate each safety function of the vehicle, from brakes and steering to physical vehicle controls like locks, headlights, and doors. With these tests, we can ensure that the vehicle operates safely in manual mode, self-driving mode with a test driver at the wheel, and fully self-driving mode without a person inside the vehicle. Overall, this testing seeks to ensure that our vehicle continues to function safely after the addition of our self-driving system.

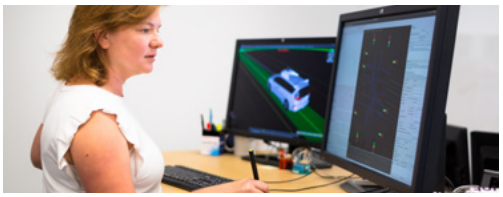


Perceptive Exhibit 2047
IPR2025-01576

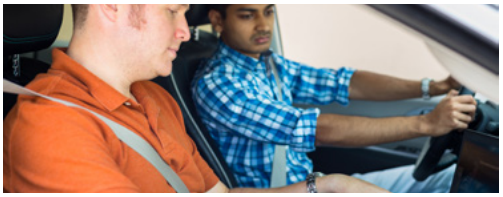
Self-Driving Software Testing

Like our hardware, our self-driving software is guided by our *Safety by Design* philosophy. We constantly and rigorously test the individual components of the software—including perception, behavior prediction, and planner—as well as the software as a whole.

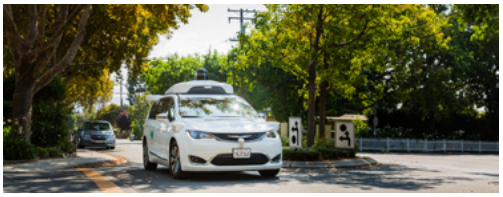
Our technology is constantly learning and improving. Each change of our software undergoes a rigorous release process. We update our software regularly for different operational design domains. Each update is tested through a combination of simulation testing, closed course testing, and driving on public roadways:



Simulation Testing:
In simulation, we rigorously test any changes or updates to our software before they're deployed in our fleet. We identify the most challenging situations our vehicles have encountered on public roads, and turn them into virtual scenarios for our self-driving software to practice in simulation.



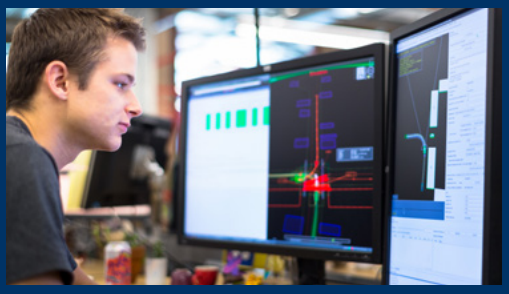
Closed-Course Testing:
New software is pushed to a few vehicles first so that our most experienced drivers can test the new software on our private test track. We can use different releases of software for different vehicles so that we can test new or specific features within different operational design domains.



Real-World Driving:
Once we confirm that our software is working as intended, we begin introducing the new software to our vehicles on public roads. We start small—our self-driving vehicles must show they can safely and consistently travel a predetermined route—and then we push the software update to our entire fleet. The more miles we travel on public roads, the more opportunities to monitor and assess the performance of software.

As we drive more road miles, we continue to further refine our driving and update our software. This continual feedback loop allows us to build confidence that our software reacts and responds appropriately in the operational design domain, enabling our vehicle to operate at SAE Level 4 safely.

Perceptive Exhibit 2047
IPR2025-01576





Waymo engineers build virtual scenarios that allow our self-driving vehicles to drive up to 8 million simulated miles each day.

Simulation: How the Virtual World Helps Our Cars Learn Advanced Real-World Driving Skills

Waymo's simulator can replay the real-world miles we have driven with each new software version, but also can build completely new realistic virtual scenarios for our software to be tested against. Each day, as many as 25,000 virtual Waymo self-driving vehicles drive up to eight million miles in simulation, refining old skills and testing out new maneuvers that help them navigate the real world safely.

For example: at the corner of South Longmore Street and West Southern Avenue in Mesa, Arizona, there's a flashing yellow arrow for left turns. This type of intersection can be tricky for humans and self-driving vehicles alike—drivers must move into a five-lane intersection and then find a gap in oncoming traffic. A left turn made too early may pose a hazard for oncoming traffic; a turn made too late may frustrate drivers behind.

Simulation lets us turn a single real-world encounter like this into thousands of opportunities to practice and master a skill.



Our self-driving vehicle encounters a flashing yellow left turn arrow in Mesa, Arizona.

Perceptive Exhibit 2047
IPR2025-01576

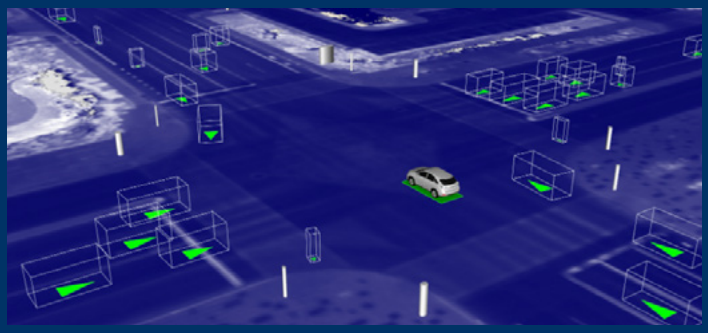
How Simulation Works



We can recreate a highly-detailed, realistic virtual version of the East Valley.

Step 1: Start with a Highly-Detailed Vision of the World

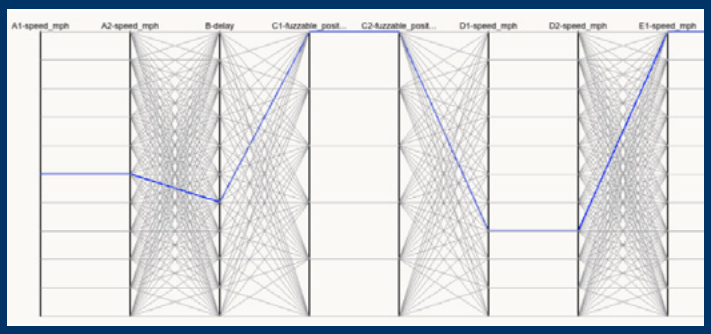
Using a powerful suite of custom-built sensors, we build a virtual replica of the intersection, complete with identical dimensions, lanes, curbs, and traffic lights. In simulation, we can focus on the most challenging interactions—flashing yellow signals, wrong-way drivers, or nimble cyclists—rather than on monotonous highway miles.



In simulation, we can practice driving the same intersection, in the same driving conditions thousands of times, with different vehicles from our fleet. In this image, we're simulating driving the intersection with one of our Lexus vehicles.

Step 2: Drive, Drive, and Redrive

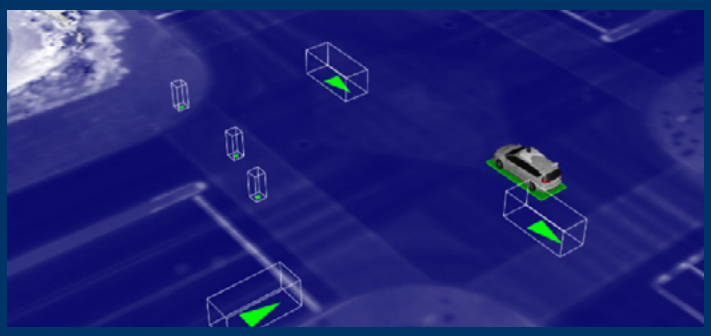
With this flashing yellow left turn now digitized in our virtual world, our software can practice this scenario thousands of times over. Every time we update the software, we can test the change at the same intersection in a variety of driving conditions. That's how we were able to teach our vehicles to naturally inch forward at that flashing yellow light, and slot in after oncoming traffic. What's more, in simulation we can practice this new skill on every flashing yellow arrow we have ever come across, in order to improve our software even faster.



Through a process called fuzzing, we can alter the speed, trajectory and position of objects on these virtual streets.

Step 3: Create Thousands of Variations

Next, we can multiply this one tricky left turn to explore thousands of variable scenarios and "what ifs?" Through a process called fuzzing, we alter the speed of oncoming vehicles and the timing of traffic lights to make sure our vehicles can still find a safe gap in traffic. The scene can be made busier and more complex by adding simulated pedestrians, motorcycles "splitting the lane," or even joggers zig-zagging across the street—all to see how that might change our driving.



To make a scene more complex, we can add vehicles, pedestrians and cyclists that never existed in the original scene.

Step 4: Validate and Iterate

Success: Our self-driving vehicle has learned how to turn confidently at a flashing yellow arrow. That new skill becomes part of our permanent knowledge base, shared with every vehicle across the fleet. In turn, we'll use real-world driving and our private closed course testing facility to validate our simulated experience. And then the cycle begins again. Each of these eventful simulator miles is guiding us to what everyone wants: billions of safe and uneventful miles in the real world.

Perceptive Exhibit 2047
IPR2025-01576



Field Tests at Our Closed-Course Facility

Waymo has set up a private, 91-acre, closed-course testing facility in California specially designed and built for our own unique testing needs. This private facility, nicknamed "Castle," is set up like a mock city, including everything from high-speed roads to suburban driveways to a railroad crossing. Our team uses this and other closed-course facilities to validate new software before it's released to our fleet of vehicles on the road, and also to stage challenging or rare scenarios so our vehicles gain experience with unusual situations.

On our closed course, we're able to conduct thousands of "structured tests" which recreate specific scenarios for learning and testing. To power our simulator, we've developed more than 20,000 simulation scenarios at Castle. Each recreates a driving situation we want to practice—an aggressive driver barreling out of a driveway, or a pedestrian suddenly emerging from a parked car—that might take hundreds of thousands of driving miles to encounter on public roads. We've staged people jumping out of canvas bags or porta potties on the side of the road, skateboarders lying on their boards, and thrown stacks of paper in front of our sensors. This "structured testing" is key to accelerating the progress of our technology and ensuring safety of our vehicles in both everyday and challenging driving situations.

Behavioral Competencies for Normal Driving

A fully self-driving vehicle must be able to handle all the everyday driving tasks expected of human drivers within the same operational design domain. This means self-driving systems need to demonstrate they have the adequate skills—or "behavioral competencies"—required for the intended locations and conditions of operation.

The U.S. Department of Transportation has recommended that Level 3, Level 4, and Level 5 self-driving vehicles should be able to demonstrate at least 28 core competencies adapted from research by California Partners for Advanced Transportation Technology (PATH) at the Institute of Transportation Studies at University of California, Berkeley. DOT also encourages companies "to consider all known behavioral competencies in the design, test, and validation" of a self-driving system.

Waymo's safety program has expanded the 28 core competencies in both breadth and depth, for which we test thousands of scenario variations—ranging in complexity—ensuring that our system can safely handle the challenges of real-world environments. In addition, we have identified further categories that expand upon the initial 28 core competencies. [17] (For a subset of Waymo's behavioral competencies, see Appendix A.)

For each competency, Waymo's team creates a wide variety of individual tests to run on our closed course facility and in simulation. For example, to test our ability to make unprotected left turns, we stage dozens of real-life situations and test to see if our vehicles respond appropriately. We include challenging variations of this common road maneuver, including using multiple lanes of oncoming traffic, obstructing our vehicle's field of view with a large truck, or providing a short green traffic light to make the turn.

For each of these scenarios we then use our simulator to create hundreds of different variations of the same encounter. With our virtual world testing, we can also create entirely new scenarios of unprotected left-hand turns so we can test this skill further. As we expand our operational design domain, the number of core competencies may grow (for example, to drive in northern U.S. states year-round, our system must be able to safely drive in snow) and the number of tests within each category may expand with more unique or complex scenarios.

While this type of scenario testing can demonstrate our software's core driving skills, these competencies need to translate out into the real world. That's why this acts merely as a starting point: our validation then moves onto testing our vehicle, hardware, and software as an integrated fully self-driving vehicle on public roads, where it demonstrates these competencies daily in real traffic situations.

Perceptive Exhibit 2047
IPR2025-01576



Testing the Fully Integrated Self-Driving Vehicle

After testing the base vehicle, the self-driving system, and the software individually, we then test the fully integrated self-driving vehicle. This includes closed-course collision avoidance testing, reliability and durability testing, and on-road testing with trained test drivers at the wheel.

Testing on Public Roads

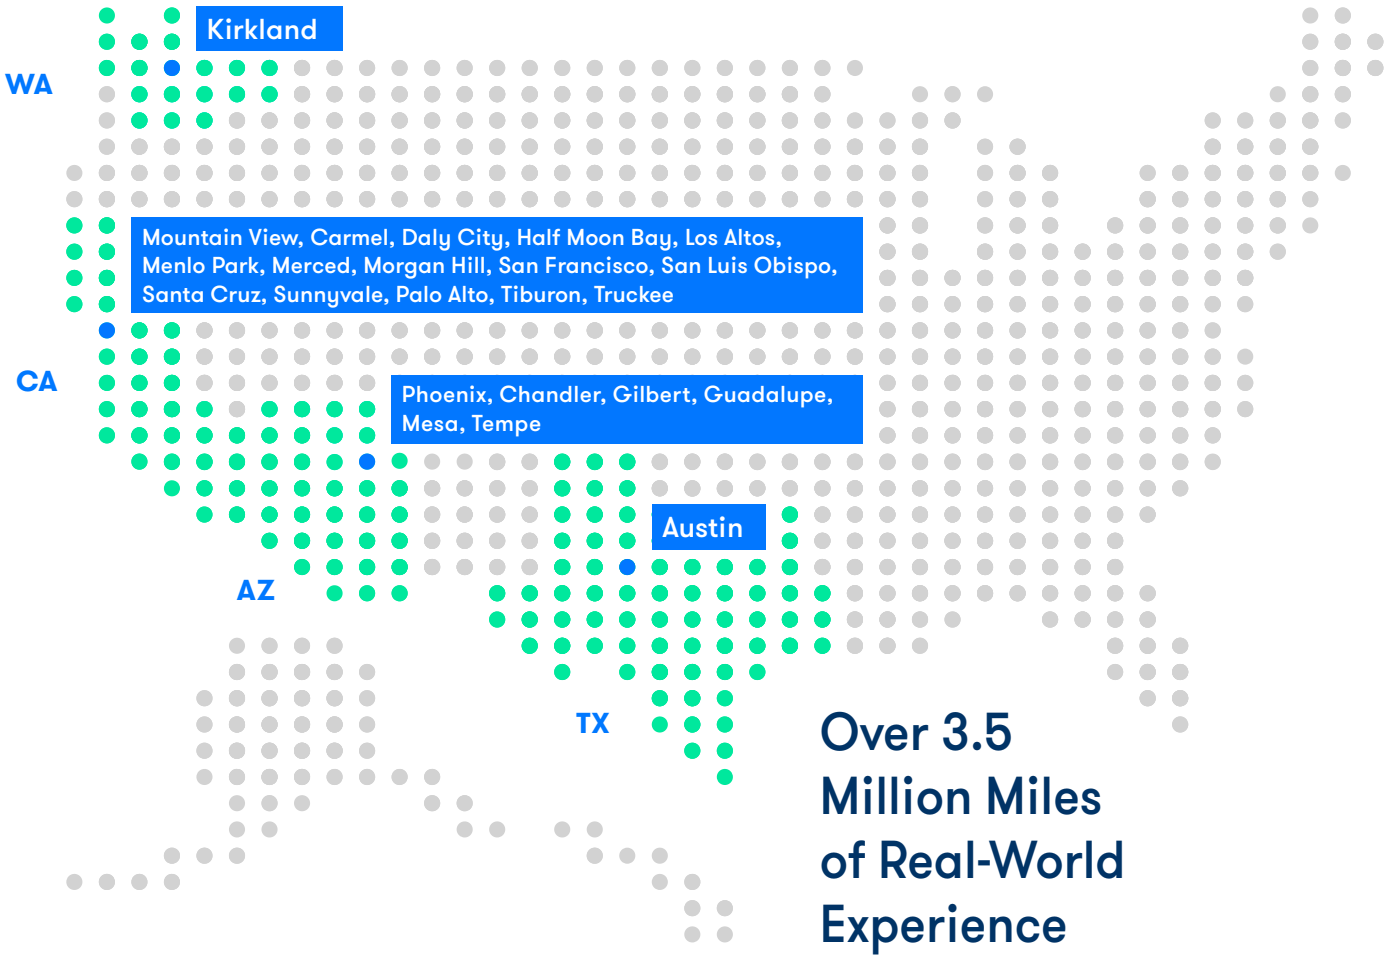
Waymo has a comprehensive on-road testing program that has been improved and refined continuously over our eight-year history. It's a critical step that allows us to validate the skills we have developed, uncover new challenging situations, and develop new capabilities.

The safety of our on-road testing program begins with highly-trained drivers. Our test drivers undergo extensive classroom training, learning about the overall system and how to monitor the vehicle safely on public roads, including taking defensive driving courses. After this training, our drivers are responsible for monitoring the system and if needed, taking control of the vehicle while we test on public roads.

Our on-road testing program drives tens of thousands of weekly miles that are used to evaluate our software. We monitor our systems to ensure they demonstrate our behavioral competencies, and we look for situations where we can build on these competencies and enable smoother driving.

Real-world testing provides a continuous feedback loop that lets us refine our system continually. Our engineers observe real-world situations, make adjustments to the software to refine our driving, and then implement those changes. This iterative approach to testing and public-road validation helps us safely scale our capabilities as we expand our operational design domain and the capabilities of our vehicles.

Perceptive Exhibit 2047
IPR2025-01576



Over 3.5 Million Miles of Real-World Experience



Real-World Experience

Over the last eight years, Waymo has tested our vehicles in four U.S. states and self-driven in more than 20 cities—from sunny Phoenix, AZ to rainy Kirkland, WA—accumulating more than 3.5 million autonomous miles in the process. As we expand to new locations, we're able to gather a wider variety of experience with different road environments, streetscapes, and driver habits.

For example, driving in Phoenix has allowed us to test our sensors and software in desert conditions, including extreme temperatures and dust in the air. We learned how to navigate more confidently around new types of vehicles, like watering trucks that move 3 mph on 45 mph roads while spraying plants in road medians. Austin provided horizontal traffic signals for the first time, while Kirkland gave us more wet weather practice.

In every new city, we meet people who aren't used to seeing self-driving cars every day. That lets us also hear fresh perspectives from diverse populations—how people want to use self-driving vehicles, what they think of our driving, and more—who together inform how we develop and refine our self-driving technology.

Perceptive Exhibit 2047
IPR2025-01576



Self-Driving in Extreme Temperatures

Our self-driving vehicles need to operate reliably and safely in extreme cold and oppressive heat. Waymo engineers have developed both our self-driving hardware and software in-house to create a complete system that can work reliably in the toughest environments.

Heat poses challenges for all modern technology. Everyday electronics like cell phones can overheat and switch off when used in the bright sun. However, our self-driving system needs to operate safely even in hot conditions. Our cars are equipped with a special cooling system that lets them operate under very hot temperatures, even with an engine running at full power and systems at full capacity.

Our engineers perform extensive testing in a wind tunnel that can mimic almost any weather condition, including the hottest temperatures ever recorded on Earth.

In addition to wind tunnel testing, we have tested our self-driving vehicle in three of America's hottest places: Las Vegas, Davis Dam, and Death Valley. The Davis Dam, on the Arizona and Nevada border, has long stretches of steep desert road for us to drive under the hot sun. The Las Vegas Strip lets us test our systems in countless busy lanes of stop-and-go city traffic under intense heat. Death Valley holds the record for the highest officially recorded temperature on Earth of 134°F.

During testing we closely monitor our systems, taking over 200 different measurements per second to confirm that our in-house sensor suite and compute keeps working as intended.



Testing Crash Avoidance Capabilities

In addition to testing core behavioral competencies, our engineers also conduct crash avoidance testing across a variety of scenarios. (To view a subset of Waymo's crash avoidance test scenarios, see Appendix B.)

Waymo has completed thousands of crash avoidance tests at our private test track. Each of these individual tests recreates a distinct driving scenario and allows us to analyze our vehicles' response. We then use our simulator to test these scenarios further and improve our overall software capabilities.

We draw from a variety of sources to learn which collisions to test against. They include our own analysis of sources such as NHTSA's fatal crash data base, and use of our extensive experience operating self-driving vehicles to expand on NHTSA's 37 pre-crash scenarios. We also test situations in which other road users create potentially dangerous situations, such as vehicles suddenly pulling out of driveways, large vehicles cutting across target lanes, motorcyclists weaving through traffic, and pedestrians jaywalking.

In 2015, NHTSA published data showing the distribution of the most common pre-crash scenarios. For example, just four crash categories accounted for 84% of all crashes: rear end crashes, vehicles turning or crossing at an intersection, vehicles running off the edge of the road, and vehicles changing lanes. Therefore, avoiding or mitigating those kinds of crashes is an important goal for our testing program. [18]



Crash avoidance testing at Waymo's closed course testing facility, Castle.

Perceptive Exhibit 2047
IPR2025-01576

Hardware Reliability and Durability Testing

Self-driving vehicles, like their conventional counterparts, must operate reliably. That means the vehicle and each of its individual components must function under extreme environmental conditions and over the lifetime of the vehicle.

Waymo engineers design unique stress tests. Using our knowledge of the physics of failure to accelerate environmental stresses on our vehicle and its individual components, we compress years of real-world use into days and weeks of testing.

We blast our components with ultraviolet radiation, bombard them with powerful water jets, dunk them into nearly freezing vats of water, corrode them in chambers full of salty mist, shake and shock them with powerful vibrations, and heat and freeze them for weeks at a time in temperature and humidity chambers. We analyze any failures and make design improvements to increase the reliability of our components. We monitor the health of each sensor, and the vehicle itself, so we can identify and fix potential failures before they occur.



Perceptive Exhibit 2047
IPR2025-01576

SECTION 4



Interacting Safely with the Public

Waymo's Early Rider Program

We want to learn how a self-driving vehicle could fit into people's every day transportation needs—whether that's as a personal use vehicle, as a ride-share, or to make public transit more accessible. That's why in April 2017, we launched our early rider program, the first public trial of our self-driving vehicles, in the Phoenix metropolitan area.

Our riders come from all walks of life, from families with teenage kids to young professionals. They're using our vehicles for everyday activities—from commuting to work to taking the kids to soccer practice. Educating our early riders on how to use our vehicles is critical. Our research team works with each new rider to provide them with information about the program and how to use our vehicles, and also how to provide feedback.

For the last 100 years, vehicles have been designed with a human driver in mind. The experiences of our early riders will teach us about how people want to interact with our vehicles, and what it's like to ride as a passenger instead of a driver. Their experiences will help us create an in-car experience that is even more intuitive and easy-to-use.

Our vehicles are designed to drive themselves, so our user interface focuses on passengers, not drivers.

That's why we've developed specific in-car features and user interfaces that help our passengers understand what our vehicles are doing on the road and let them do things like set a destination, ask the vehicle to pull over, and get in touch with Waymo support staff as needed. We also understand the transportation challenges that exist today, especially around accessibility, and we are working to develop solutions that work for riders of all abilities.

In addition to creating a safe and intuitive everyday ride for our passengers, Waymo has also developed procedures in case of emergency. For example, not only are our vehicles designed to detect collisions and respond appropriately to emergency vehicles on the road, but we have also conducted trainings with law enforcement and first responders who may come into contact with our vehicles.

Finally, the potential of self-driving cars will only be realized by growing public awareness and acceptance of this technology. In October 2017, Waymo helped to launch *Let's Talk Self-Driving* (letstalkselfdriving.com), the world's first public education campaign about fully self-driving vehicles. Working in partnership with national and local safety, mobility, and seniors groups, the initiative hopes to engage and educate the public about how this technology works and the enormous benefits self-driving technology could unlock.

Perceptive Exhibit 2047
IPR2025-01576

Making Waymo's Vehicles Easy to Use



Display

The Waymo passenger display screen shows important trip information such as destination and time to arrival. It also visualizes static road elements like traffic lights, stop signs, and dynamic agents in the environment such as vehicles, cyclists, and pedestrians. That way, riders can understand what the vehicle is perceiving and responding to, and be confident in the vehicle's capabilities.



Start Ride Button

Riders can start the ride whenever they're ready, using the mobile app or a button inside the vehicle.



Pull Over Button

The vehicle features a "Pull Over" button for its riders. When pressed, the vehicle will identify the nearest location to safely pull over so that the rider can exit the vehicle before their original destination.



Mobile App

Participants in Waymo's early rider program use a mobile app to request a ride in a Waymo vehicle to their intended destination. The app also allows users to give ride feedback and contact Waymo's rider support.



Rider Support Team

Waymo has created a rider support team to help answer questions for our early riders. These specialists can be reached with a button-press inside the vehicle or by calling or chatting with our rider support team from the mobile app. Our rider support specialists can speak with riders during the regular course of a trip or assist in case of an emergency.

Rider Experience

Waymo's user experience is guided by four main principles: give passengers the information they need for a seamless trip; help passengers anticipate what's next; proactively communicate the vehicle's response to events on the road; and help passengers engage safely with the vehicle.

Audio and visual information provided to passengers helps them know what to expect, reminds them of safety features such as seat belts, and permits them to communicate with Waymo's rider support personnel.

We also want our passengers to be aware of what the vehicle is perceiving, and why it is taking specific actions. Each vehicle also provides occupants with useful visual and audio information throughout the trip, to help them understand what the vehicle and other road users around it are doing. In Waymo's self-driving minivans, the in-vehicle screens are used to provide visual ride information, such as destination, current speed, and the route the vehicle intends to take. An audio system provides audible notifications and cues to all riders.

In the event of a safety-critical event, the screens and audio system are designed to provide the occupant with specific visual and audio cues depending on the nature of the event. We've designed multiple ways for our riders to interact with our vehicle, whether it's through the pressing of physical buttons, a mobile app, or by speaking with a Waymo rider support specialist.

Perceptive Exhibit 2047
IPR2025-01576

Accessibility Features In Development

An Accessible Mobile App:

We're building our mobile app to be intuitive and accessible. It's designed for use with Android TalkBack, iOS VoiceOver, and other accessibility services.

Audio Cues and Tools:

Visually impaired riders may need help locating our vehicles at their pickup locations. We're exploring specific "wayfinding" features, including ways that these riders can ask their vehicle to make a sound to help guide them to the vehicle. Additional audio cues can be turned on in the mobile app and will be available in the vehicle to keep the rider informed of their journey.

Braille Labels:

The ride buttons in our self-driving vehicles are accompanied by Braille to allow vision-impaired riders to start the ride, pull over the vehicle, or call to speak to an operator who can provide further assistance and information. These buttons are also available in the mobile app.

Visual Display:

Through every phase of the ride, deaf and hearing-impaired riders will have access to on-screen visual cues of what is happening around the vehicle.

Accessible Rider Support:

Our chat-based rider support will be available to all riders of all abilities through visual displays or audio inside the vehicle.

Accessibility: Unlocking Opportunities for Those Who Cannot Drive Today

We believe our technology holds the potential to improve safety and mobility for people around the world. From the start, Waymo has been listening to and working with the disability community. We continue to learn about the unique needs of different riders, and what we learn will inform new features that will make the experience accessible to people who have historically had to rely on others to get around.

We also know we can't reach our goals alone. Waymo is committed to working with our partners to identify vehicle platforms and solutions that can serve a broader set of individuals.



Perceptive Exhibit 2047
IPR2025-01576



Testing with the Chandler Arizona Police Department

We've collaborated with the Chandler Police and Fire departments in Arizona to conduct emergency vehicle testing with our self-driving minivans. Our powerful suite of sensors, including our long-range audio detection system, observed local police vehicles, motorcycles, ambulances, firetrucks, and undercover vehicles as they trailed, passed, and led our vehicles. Our sensors collected samples at various speeds, distances, and angles—building up a library of sights and sounds that will help our vehicles respond safely to emergency vehicles on the road.

Emergencies and Interacting With Law Enforcement and First Responders

Our self-driving vehicles are designed to interact with law enforcement and first responders safely on road. Using our suite of custom-built sensors, including an audio detection system, our software can identify a nearby fire truck, detect its flashing lights, and hear sirens up to hundreds of feet away. Our audio sensors are designed to discern the direction sirens are likely coming from, improving our vehicles' ability to respond in both a safe and timely manner. Once an emergency vehicle is detected, our vehicle can respond by yielding, pulling over to the side of the road, or coming to a complete stop.

Waymo also briefs local authorities in every city in which we test, and offers a line of communication for further engagement. In some cities, Waymo has also conducted on-site training to help police and other emergency workers identify and access our vehicle in emergency situations. We plan to continue conducting these on-site trainings, while expanding the scope of the training program as our vehicles become more capable and our operational design domain expands.

Perceptive Exhibit 2047
IPR2025-01576



SECTION 5

Conclusion

For more than eight years, Waymo has focused on one thing: bringing fully self-driving technology to the world. We are committed to *Safety By Design*, and we have built a culture that puts safety, and open communication about safety, at its core. All of us at Waymo are committed to the goal of making it safe and easy for everyone to get around.

This report summarizes our efforts to ensure the safe deployment of fully self-driving vehicles that use Waymo technology. We are excited about the potential this technology holds to improve road safety and provide new mobility options for the world. For further information about Waymo's self-driving technology, please visit www.waymo.com.

Perceptive Exhibit 2047
IPR2025-01576

APPENDIX

Scenario Types Used for Testing and Validation

Waymo tests our vehicles comprehensively to ensure that they are capable of operating safely in reasonably foreseeable scenarios that could present a safety hazard.

The following types of scenarios are illustrative of the breadth of our testing program and are designed to ensure our vehicles have: 1) basic behavioral competencies and 2) the ability to avoid or mitigate crashes in common crash scenarios.

Perceptive Exhibit 2047
IPR2025-01576

Appendix A. Basic Behavioral Competency Testing

We believe that our fully self-driving vehicles should be able to successfully demonstrate competency in a variety of reasonably foreseeable traffic situations that are within the vehicle's operational design domain. Our system can recognize and stay within its design domain, and the set of competencies expands or shrinks in accordance with the scope of each operational design domain. For each behavioral competency shown in the table below, we test a wide range of scenarios with variations in factors such as road configuration, the speed of our vehicle or other vehicles, and lighting conditions.

Set of Behavioral Competencies Recommended by NHTSA	
1	Detect and Respond to Speed Limit Changes and Speed Advisories
2	Perform High-Speed Merge (e.g., Freeway)
3	Perform Low-Speed Merge
4	Move Out of the Travel Lane and Park (e.g., to the Shoulder for Minimal Risk)
5	Detect and Respond to Encroaching Oncoming Vehicles
6	Detect Passing and No Passing Zones and Perform Passing Maneuvers
7	Perform Car Following (Including Stop and Go)
8	Detect and Respond to Stopped Vehicles
9	Detect and Respond to Lane Changes
10	Detect and Respond to Static Obstacles in the Path of the Vehicle
11	Detect Traffic Signals and Stop/Yield Signs
12	Respond to Traffic Signals and Stop/Yield Signs
13	Navigate Intersections and Perform Turns
14	Navigate Roundabouts
15	Navigate a Parking Lot and Locate Spaces
16	Detect and Respond to Access Restrictions (One-Way, No Turn, Ramps, etc.)
17	Detect and Respond to Work Zones and People Directing Traffic in Unplanned or Planned Events
18	Make Appropriate Right-of-Way Decisions
19	Follow Local and State Driving Laws

Perceptive Exhibit 2047
IPR2025-01576

20	Follow Police/First Responder Controlling Traffic (Overriding or Acting as Traffic Control Device)
21	Follow Construction Zone Workers Controlling Traffic Patterns (Slow/Stop Sign Holders)
22	Respond to Citizens Directing Traffic After a Crash
23	Detect and Respond to Temporary Traffic Control Devices
24	Detect and Respond to Emergency Vehicles
25	Yield for Law Enforcement, EMT, Fire, and Other Emergency Vehicles at Intersections, Junctions, and Other Traffic Controlled Situations
26	Yield to Pedestrians and Bicyclists at Intersections and Crosswalks
27	Provide Safe Distance From Vehicles, Pedestrians, Bicyclists on Side of the Road
28	Detect/Respond to Detours and/or Other Temporary Changes in Traffic Patterns
Examples of Additional Behavioral Competencies Tested by Waymo	
29	Moving to a Minimum Risk Condition When Exiting the Travel Lane is Not Possible
30	Perform Lane Changes
31	Detect and Respond to Lead Vehicle
32	Detect and Respond to a Merging Vehicle
33	Detect and Respond to Pedestrians in Road (Not Walking Through Intersection or Crosswalk)
34	Provide Safe Distance from Bicyclists Traveling on Road (With or Without Bike Lane)
35	Detect and Respond to Animals
36	Detect and Respond to Motorcyclists
37	Detect and Respond to School Buses
38	Navigate Around Unexpected Road Closures (e.g. Lane, Intersection, etc.)
39	Navigate Railroad Crossings
40	Make Appropriate Reversing Maneuvers
41	Detect and Respond to Vehicle Control Loss (e.g. reduced road friction)
42	Detect and Respond to Conditions Involving Vehicle, System, or Component-Level Failures or Faults (e.g. power failure, sensing failure, sensing obstruction, computing failure, fault handling or response)
43	Detect and Respond to Unanticipated Weather or Lighting Conditions Outside of Vehicle's Capability (e.g. rainstorm)
44	Detect and Respond to Unanticipated Lighting Conditions (e.g. power outages)
45	Detect and Respond to Non-Collision Safety Situations (e.g. vehicle doors ajar)
46	Detect and Respond to Faded or Missing Roadway Markings or Signage
47	Detect and Respond to Vehicles Parking in the Roadway

Perceptive Exhibit 2047
IPR2025-01576

Appendix B. Avoidance or Mitigation of Common Crash Scenarios

Certain types of crashes account for a substantial percentage of all crashes. Avoiding or mitigating those kinds of crashes, therefore, is an important goal for our vehicle development program. In late 2015, NHTSA published data showing the distribution of pre-crash scenarios. [19]

Four scenarios accounted for the vast majority of crashes:

- 29 percent of the vehicles were involved in rear-end crashes
- 24 percent of the vehicles were turning or crossing at intersections just prior to the crashes
- 19 percent of the vehicles ran off the edge of the road
- 12 percent involved vehicles changing lanes

Therefore, these scenarios figure prominently in the evaluation of our vehicles. The table below illustrates just a few of the test scenarios we employ to determine our vehicle's ability to avoid or mitigate crashes in these all-important situations, as well as in other crash situations.

Crash Avoidance Category	Example Test Scenario
Rear-end Demonstrate ability to avoid or mitigate crashes with lead vehicles.	Fully self-driving vehicle approaches stopped lead vehicle
	Fully self-driving vehicle approaches disabled vehicle
	Fully self-driving vehicle approaches lead vehicle traveling at lower constant speed
	Fully self-driving vehicle approaches lead vehicle traveling at slower speed and initiating strong braking
	Fully self-driving vehicle approaches lead vehicle accelerating
	Fully self-driving vehicle following a lead vehicle making a maneuver (e.g. cutting into lane or pulling out of driveway)
	Fully self-driving vehicle approaches lead vehicle decelerating
	Fully self-driving vehicle approaches other vehicle(s) reversing
	Fully self-driving vehicle approaches other vehicle(s) parking
Intersection Demonstrate ability to detect vehicle entering path at perpendicular angle and apply brakes.	Fully self-driving vehicle approaches protected intersection, Vehicle A approaches from right
	Fully self-driving vehicle approaches protected intersection, Vehicle A approaches from left
	Fully self-driving vehicle prepares to turn across unprotected intersection, oncoming Vehicle A approaches
	Crossing path collisions - other vehicle running red light
	Crossing path collisions - other vehicle running stop sign

Perceptive Exhibit 2047
IPR2025-01576

Crash Avoidance Category (continued)	**Example Test Scenario** (continued)
Road Departure Demonstrate ability to steer clear of roadway edge and stay within lane.	Fully self-driving vehicle travels down straight road (with or without prior vehicle maneuver)
	Fully self-driving vehicle travels down curved road (with or without prior vehicle maneuver)
	Fully self-driving vehicle travels down straight road with visible lane marking
	Fully self-driving vehicle travels down straight road with faded or missing lane marking
	Fully self-driving vehicle travels down curved road with visible lane marking
	Fully self-driving vehicle travels down curved road with faded or missing lane marking
	Fully self-driving vehicle travels down wet road with lane marking
	Fully self-driving vehicle approaches other vehicle(s) reversing
	Fully self-driving vehicle travels down wet road with faded or missing lane marking
Lane Change Demonstrate ability to avoid or mitigate crash when other vehicles make lane changes or merge.	Lane changes - other vehicles turning same direction
	Lane changes - other vehicles parking same direction
	Lane changes - other vehicles changing lanes same direction
	Lane changes - other vehicles drifting same direction
	Lane merges

However, we evaluate those capabilities in many more situations than those shown here. We have developed many additional test scenarios based on NHTSA's overall pre-crash scenarios, our analysis of additional sources such as NHTSA's fatal crash data base, and from our own extensive experience operating self-driving vehicles. [19]

Perceptive Exhibit 2047
IPR2025-01576

GLOSSARY

Behavioral Safety. An aspect of system safety that focuses on how a system should behave normally in its environment to avoid hazards and reduce the risk of mishaps: for instance, detect objects and respond in a safe way (slow down, stop, turn, lane change, etc.).

California Partners for Advanced Transportation Technology (PATH). A research and development program of the University of California, Berkeley, with staff, faculty, and students from universities worldwide and cooperative projects with private industry, state, and local agencies, and nonprofit institutions. See www.path.berkeley.edu.

Crash Safety. An aspect of system safety that focuses on reducing the consequences of collisions by reducing the severity of the event as experienced by vehicle occupants or other road users.

Dynamic Driving Task. All of the real-time operational and tactical functions required to operate a vehicle in on-road traffic, excluding strategic functions such as trip scheduling and selection of destinations and waypoints.

Fault. An abnormal condition in the system. A fault might be triggered by hardware failures, software error detection, detection of off-nominal system performance, or other conditions defined within the diagnostics capability of the system.

Functional Safety. An aspect of system safety that focuses on how the system should detect and respond to failures, errors, or off-nominal performance of the self-driving system (e.g., fail operational, fail safe, or transition to a minimal risk condition).

Hazard. Any real or potential condition that can cause injury, illness, or death to personnel; damage to or loss of a system, equipment or property; or damage to the environment. (MIL-STD-882E).

Hazard Analysis. A process of identifying or recognizing hazards that may arise from a system or its environment, and analyzing their potential causes for the purpose of assessing risk and initiating actions necessary to reduce the risk to acceptable levels. Results of hazard analyses are also used to develop verification and validation approaches and procedures to demonstrate that hazard risks have been mitigated to acceptable levels.

Manufacturer. An individual or company that manufactures self-driving vehicles or equipment for testing and deployment on public roadways.

Minimal Risk Condition. A low-risk operating mode in which a fully self-driving vehicle operating without a human driver achieves a reasonably safe state, such as bringing the vehicle to a complete stop, upon experiencing a failure of the self-driving system that renders the vehicle unable to perform the entire dynamic driving task.

Mishap. An event or series of events resulting in death, injury, illness, or damage to property.

Mishap Risk. See Definition for Risk.

Object and Event Detection and Response. The perception by the system of any circumstance that is relevant to the immediate driving task, as well as the appropriate driver or system response to such a circumstance.

Operational Design Domain. A description of the specific operating conditions in which a self-driving system is designed to properly operate, including but not limited to roadway types, speed range, environmental conditions (weather, daytime/nighttime, etc.), and other domain constraints.

Perceptive Exhibit 2047
IPR2025-01576

Operational Safety. An aspect of system safety that focuses on the interaction between our vehicles and passengers.

Non-Collision Safety. An aspect of system safety that focuses on physical non-collision hazards.

Requirement. A general term used to describe the set of statements that identifies a system's functions, characteristics, or constraints.

Risk. An expression of the possibility and impact of a mishap in terms of hazard severity and hazard probability of occurrence. It routinely reflects conditions such as personnel error, environmental conditions, design characteristics, procedural deficiencies, or subsystem or component failure or malfunction.

SAE J3016. "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," published by SAE International in September 2016.

Safety. Freedom from those conditions that can cause death, injury, occupational illness, damage to or loss of equipment or property, or damage to the environment. (MIL-STD-882E).

Safety Requirement. 1) A system or subsystem requirement that is associated with a hazard mitigation or reduces the risk of an identified hazard. 2) A regulatory safety requirement generated from a governing agency. 3) Safety requirement derived from an industry standard or published best practice.

Self-Driving System. A Level 4 or 5 system which has hardware and software that are collectively capable of performing the entire dynamic driving task, without a human driver. This distinguishes it from Level 1, 2, or 3 systems that require a human driver.

Fully Self-Driving Vehicle. A vehicle equipped with a self-driving system designed to function without a human driver as a level 4 or 5 system.

Subsystem. 1) A grouping of items satisfying a logical group of functions within a particular system. (MIL-STD-882E) 2) A major part of a system which in itself has the characteristics of a system, usually consisting of several components.

System. 1) An integrated composite of people, products, and processes that provide a capability to satisfy a stated need or objective. (MIL-STD-882E) 2) The organization of hardware, software, material, facilities, personnel, data, and services needed to perform a designated function within a stated environment with specified results.

System Safety. 1) The application of engineering and management principles, criteria, and techniques to achieve acceptable mishap risk, within the constraints of operational effectiveness and suitability, time, and cost, throughout all phases of a system life cycle. (MIL-STD-882E). 2) The optimum degree of safety within the constraints of operational effectiveness, time, and cost attained through specific application of system safety engineering throughout all phases of a system (McGraw Hill Dictionary of Technical Terms).

System Safety Engineering. 1) An engineering discipline that employs specialized professional knowledge and skills in applying scientific and engineering principles, criteria, and techniques to identify and eliminate hazards, in order to reduce the associated mishap risk. (MIL-STD-882E) 2) An element of systems management involving the application of scientific and engineering principles for the timely identification of hazards, and initiation of those actions necessary to prevent or mitigate hazards within the system.

Perceptive Exhibit 2047
IPR2025-01576

END NOTES

[1] "Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey." National Highway Traffic Safety Administration, February 2015. https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812115

[2] "Global Status Report on Road Safety 2015." World Health Organization, 2015. http://www.who.int/violence_injury_prevention/road_safety_status/2015/en/

[3] "2016 Fatal Motor Vehicle Crashes: Overview." National Highway Traffic Safety Administration, October 2017. https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812456; "Quick Facts 2015." National Highway Traffic Safety Administration, May 2017 (updated). https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812348

[4] "The Economic and Societal Impact Of Motor Vehicle Crashes, 2010." National Highway Traffic Safety Administration, May 2014, DOT HS 812 013. https://www-nrd.nhtsa.dot.gov/Pubs/812013.pdf

[5] "The Economic and Societal Impact Of Motor Vehicle Crashes, 2010." National Highway Traffic Safety Administration, May 2015 (revised). https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812013

[6] ibid.

[7] Schrank, D., Eisele, B., Lomax, T., and Bak, J. "2015 Urban Mobility Scorecard." The Texas A&M Transportation Institute and INRIX, August 2015. https://static.tti.tamu.edu/tti.tamu.edu/documents/mobility-scorecard-2015.pdf

[8] "Quieter Cars and the Safety Of Blind Pedestrians: Phase I." National Highway Traffic Safety Administration, April 2010. https://www.nhtsa.gov/DOT/NHTSA/NVS/Crash%20Avoidance/Technical%20Publications/2010/811304rev.pdf

[9] Rosenbloom, Sandra. "The Mobility Needs of Older Americans: Implications for Transportation Reauthorization." The Brookings Institution, July 2003. https://www.brookings.edu/wp-content/uploads/2016/06/20030807_Rosenbloom.pdf

[10] "MIL-STD-882E: Standard Practice: System Safety." U.S. Department of Defense, 11 May 2012.

[11] "ISO 26262: Road Vehicles – Functional Safety." The International Organization for Standardization (ISO), 15 Nov 2011.

[12] "Description of the SHRP 2 Naturalistic Database and the Crash, Near-Crash, and Baseline Data Sets." Virginia Tech Transportation Institute (VTTI), April 2016. https://vtechworks.lib.vt.edu/bitstream/handle/10919/70850/SHRP_2_CrashNearCrashBaselineReport_4-25-16.pdf?sequence=1

[13] "SAE J3016. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles." SAE International, September 2016.

[14] "Federal Automated Vehicles Policy." National Highway Traffic Safety Administration, September 2016. (See Safety Assessment notes on page 15.) https://www.transportation.gov/AV/federal-automated-vehicles-policy-september-2016

[15] As NHTSA has noted: "A minimal risk condition will vary according to the type and extent of a given failure, but may include automatically bringing the vehicle to a safe stop, preferably outside of an active lane of traffic." (See Automated Driving Systems 2.0 on page 8.)

Perceptive Exhibit 2047
IPR2025-01576

[16] Crashes are reported consistent with state law and we cooperate with law enforcement under established legal process.

[17] "Input to NHTSA's Development of Guidelines for the Safe Deployment and Operation Of Automated Vehicle Safety Technologies." Google, Inc., May 2016. https://drive.google.com/file/d/0Byq-WVq3Y4HEcWVvVjc3TXEwOEE/view?usp=sharing

[18] "New Car Assessment Program (NCAP), 80 Fed Reg 78522 at 78552, December 16, 2015) https://www.federalregister.gov/documents/2015/12/16/2015-31323/new-car-assessment-program

[19] Our test scenarios are derived from multiple sources, including: Najm, W. G., Smith, J. D., and Yanagisawa, M. "DOT HS 810 767: Pre-Crash Scenario Typology for Crash Avoidance Research." National Highway Traffic Safety Administration, April 2007. https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/pre-crash_scenario_typology-final_pdf_version_5-2-07.pdf
Data from the NHTSA's Fatality Analysis Reporting System (FARS) database https://www.nhtsa.gov/research-data/fatality-analysis-reporting-system-fars and hazardous situations Waymo has encountered during our eight years of driving experience.

Perceptive Exhibit 2047
IPR2025-01576