Waymo Safety Report

September 2020

WAYMO

WAYMO ONE

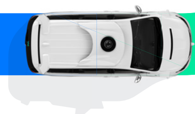

Our Mission

Waymo's mission is to bring self-driving technology to the world, making it safe and easy for people and things to get where they are going. We're building the World's Most Experienced Driver™ and we believe our technology can improve access to mobility, and save thousands of lives now lost to traffic crashes.

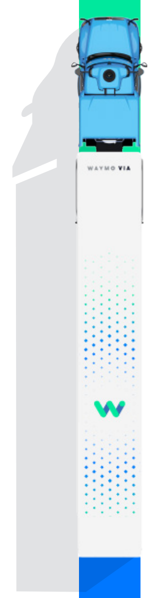

WAYMO VIA

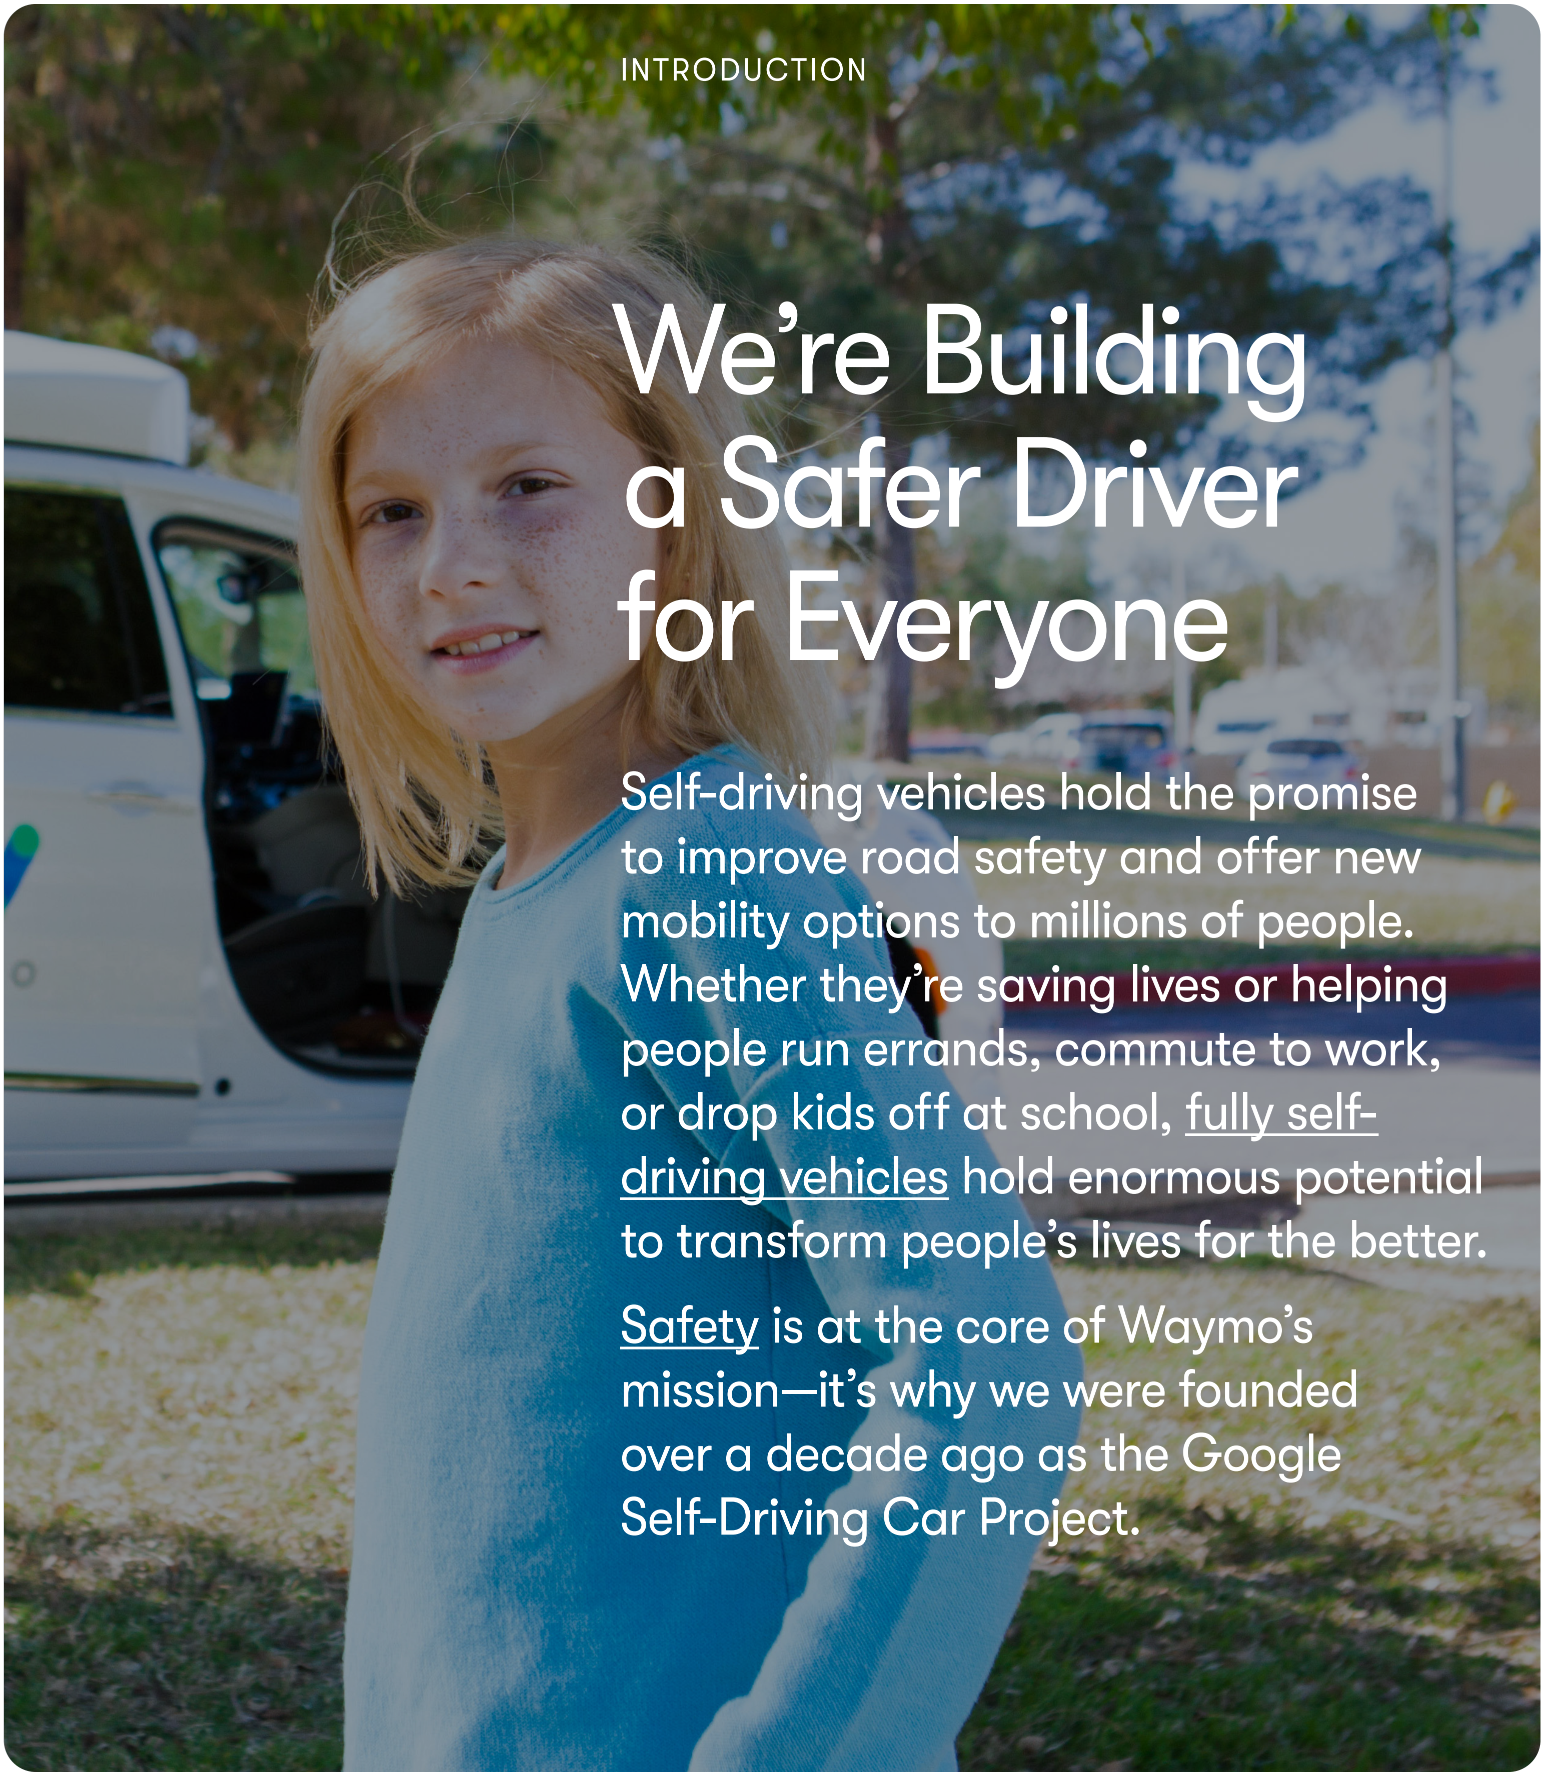

INTRODUCTION

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 a decade ago as the Google Self-Driving Car Project.

20+

Million
Self-Driven Miles

15+

Billion
Simulated Miles

Every year, 1.35 million lives are lost to traffic crashes around the world, and in the U.S. the number of tragedies is growing. 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, test and deploy our technology. In this Safety Report on Waymo's fully self-driving technology (which we call the "Waymo Driver"), we detail Waymo's work on—and our commitment to—safety. This overview of our safety program underscores the important lessons learned through the 20+ million miles Waymo's vehicles have self-driven on public roads and through our 15+ billion miles of simulated driving.

In its 2017 automated vehicle guidance, *Automated Driving Systems 2.0: A Vision for Safety*, the U.S. Department of Transportation (DOT) outlined 12 safety design elements and encouraged companies testing and deploying **self-driving systems** to address each of these areas. Over the course of this Safety Report, we address each of these and 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.

Table of Contents

The World Around Us

At Waymo, we're designing fully self-driving vehicles that make it safe and easy for people and things to get where they're going.

Self-Driving Technology Can Save Lives and Improve Mobility

Safety

1.35
MILLION
deaths worldwide due to vehicle crashes in 2016. [2]

36,560
THOUSAND
deaths in the U.S. in 2018 and 2.7 million injuries. [3]

Society

$836
BILLION
in harm from loss of life and injury each year. [4]

$242
BILLION
in annual econmoic costs. [4]

$179
BILLION
in gas burned and time lost each year. [5]

Mobility and Quality of Life

12
BILLION
people 40 years and over in the United States have vision impairment. [6]

79
PERCENT
of seniors age 65 and older live in car-dependent communities. [7]

54
HOURS
wasted in traffic each year per person. [5]

Building the World's Most Experienced Driver™

The Waymo Driver gains experience with every mile, in each car. →

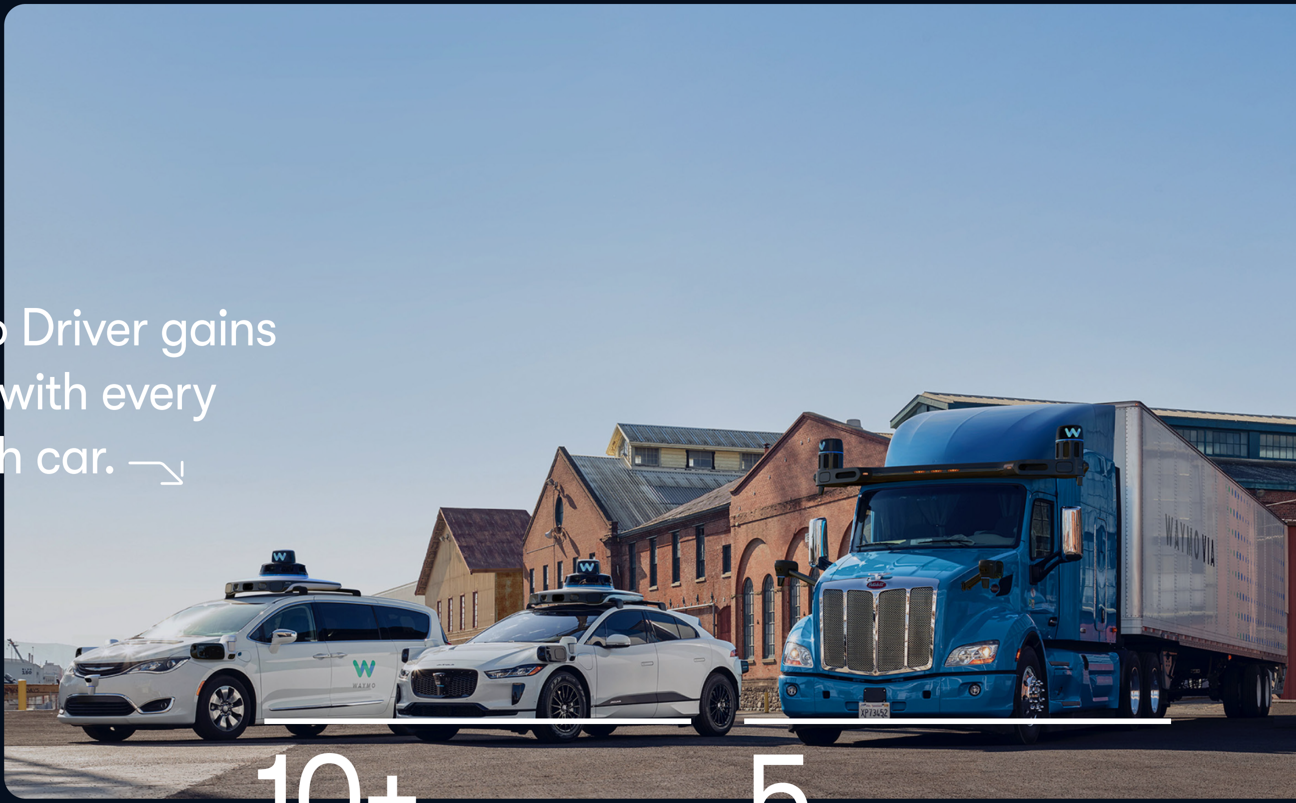

10+

More than a Decade of Self-Driving in More than 10 States

5

Generations of Self-Driving Vehicles

15+

Billion Self-Driven Miles in Simulation

20+

Million Real-World Miles on Public Roads

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.

01 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, the Waymo Driver cross-references our pre-built maps with real-time sensor data to precisely determine their location on the road.

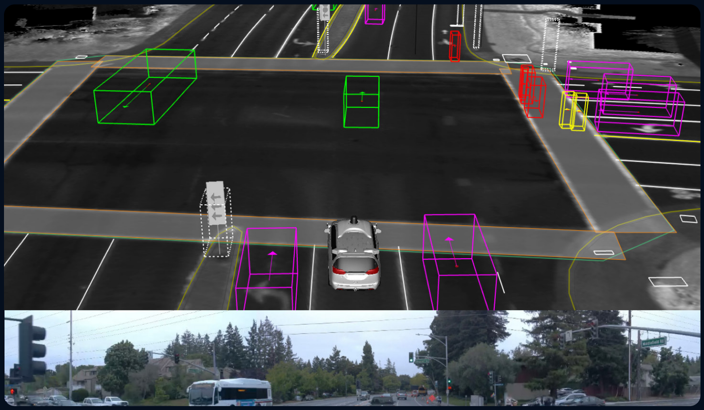

In this example, our Waymo Driver has detected vehicles (depicted by green and purple boxes), pedestrians (in yellow), and cyclists (in red) at the intersection—and a construction zone up ahead.

02 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. The Waymo sensor suite has a 360° degree view around the vehicle and is designed to respond to objects up to up to 300 meters away (nearly three football fields).

03

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, may impact the behavior of others around it.

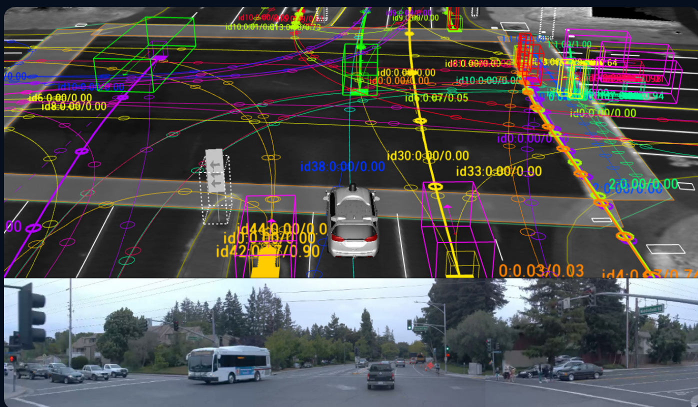

The simulated imagery shown demonstrates how Waymo's software assigns predictions to each object surrounding our vehicle—other vehicles, cyclists, pedestrians, and more.

04

What Should I Do?

Waymo's software considers all of this information as it finds an appropriate route for the vehicle to take, then selects the exact trajectory, speed, lane, and steering maneuvers needed to progress along this route safely. Because our Waymo Driver constantly monitors the environment, and predicts the future behavior of other road users in 360° degrees around our vehicles, our Driver can respond quickly and safely to any changes on the road.

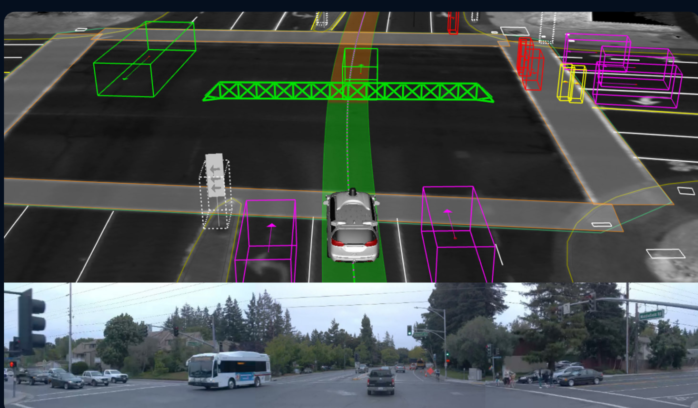

The green path indicates the trajectory through which our vehicle can proceed ahead. The series of green fences indicate that the self-driving vehicle can proceed, and that the Waymo Driver has identified the vehicles ahead and understands it has to maintain certain headway.

01 Our System Safety Program

As the first company to complete a fully self-driving trip on public roads in 2015, we have written our own playbook at Waymo.

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. Over time, we have evolved that early safety program into a comprehensive and robust *Safety by Design* approach.

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.

In line with these best practices, each individual component and all **subsystems** are tested robustly to ensure the vehicle performs safely when integrated as a complete self-driving system. In addition, we thoroughly validate that the vehicle works safely as a fully self-driving vehicle on the road, and our extensive testing program helps us 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 and fault protection, 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.

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. 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 during a crash 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 manufactuers must certify that their base vehicles meet applicable FMVSS requirements. At Waymo, we are focused on building the Driver, and we partner with automakers that share our commitment to improving road safety.

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, 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, trained drivers, first responders, or bystanders.

Safety Processes

Waymo organizes the processes we use to keep our vehicles safe through our System Safety Program. We reduce the risk of potential hazards by developing safety requirements and testing to demonstrate that safety risks have been reduced to the levels identified in the analyses and addressed in the design.

Our approach includes 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 identification methods such as customized software safety analysis, System-theoretic Process Analysis (STPA), Fault Tree Analysis (FTA), 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.

Our hazard analysis process helps identify requirements for our self-driving system's 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 through driving in simulation, on closed courses, and on public roads. We use information gathered from this testing, as well as research into national crash data and naturalistic driving studies [8], 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 allow fully self-driving operation on public roads.

SECTION

02

How Waymo's Fully Autonomous Vehicles Work

The Self-Driving System

Our fully self-driving system is designed to operate without a human driver, unlike ADAS 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, hardware, and compute that, when integrated into the vehicle, performs all driving functions.

The Waymo Driver 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 failure. Unlike driving automation 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 takes over. [9]

The Case for Full Autonomy: Allowing Passengers to Stay Passengers

Advanced driver-assist (ADAS) technologies were one of the first applications our team 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 an important reason why Waymo is committed to *fully* self-driving vehicles. Our technology takes care of all of the driving, allowing passengers to stay passengers.

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.

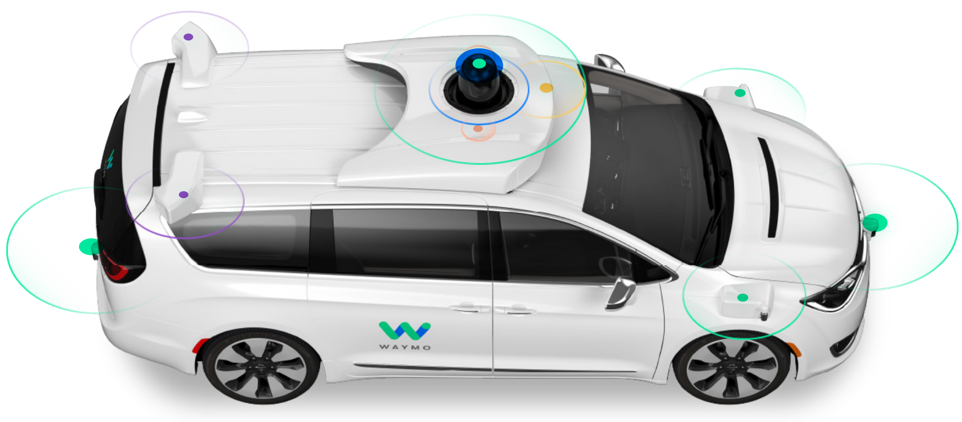

Lidar (Laser) System

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 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. 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 electromagnetic waves to perceive objects and movement. Radar remains effective in rain, fog, and snow, and operates equally well day or night. Waymo's radar system has a continuous 360° degree view, so it can track the presence and speed of road users in front, behind and to both sides of the vehicle.

Inertial Measurement Unit

This module uses accelerometers and gyroscopes with input from GPS, maps, wheels speeds, and laser and radar measurements to provide highly accurate position, velocity, and heading information to the vehicle. This information remains highly accurate even in the event of a sensor, vehicle component, or other system failure.

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.

Our Self-Driving Software

Our self-driving software is the "brain" of the Waymo Driver. 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 more than a decade 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 with over 15 billion miles of simulated driving, and more than 20 million miles of public 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 and estimates their states over time (e.g., speed, heading, and acceleration), while also producing scene understanding of the environment. 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 there's a construction zone or a lane is blocked because of the many cones in front of it.

Behavior Prediction

With behavior prediction, our software is designed to model, predict, and understand the intent of each object on the road. Because we have millions of miles of on-road 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 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.

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 operational design domain is defined by elements such as geographies, roadway types, speed range, weather, and time of day.

An operational design domain can be very limited: for instance, a single fixed route on low-speed public streets or private grounds 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 complex city streets in a variety of weather conditions and times of day within broad geographic areas.

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 Waymo Driver is designed to automatically detect sudden changes (such as a snowstorm) that would affect safe driving within its 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 our geographic areas of operation. Through our internal programs and processes, we identify legal requirements, build those requirements into our system, and evaluate our performance. Before our vehicles drive in a new jurisdiction, 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 and appropriately. For example, California and Arizona configure school zones differently (e.g. schools in Arizona place temporary signs in the roadway, while California schools don't).

Waymo's operational design domain continues to evolve. Our ultimate goal is to develop fully self-driving technology that can move people and things 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.

Our Vehicles' Redundant Safety-Critical Systems

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, the Waymo Driver is built to 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 system runs thousands of checks 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 Waymo Driver 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 our riders, our vehicle, and other road users safe, including pulling over or coming to a safe stop. [10]

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.

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.

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 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 Power Systems

Two independent power sources provide redundant power for each of the critical driving systems. These independent power sources ensure that our vehicles' critical driving functions remain online during single power failures or circuit interruptions.

Communication Systems

Self-driving system communication paths for critical vehicle control are duplicated on separate Controller Area Network (CAN) buses. This enables a safe pullover in the event of a single CAN bus or connector failure.

Redundant Collision Detection and Avoidance System

Multiple redundant 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.

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 custom maps provide our vehicles with a deep understanding of their physical environment: road types, the distance and dimensions of the road itself, and other topographical features.

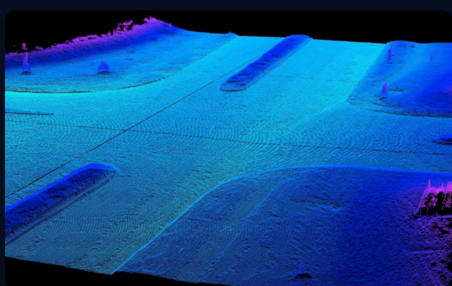

We take this data and add salient information that includes traffic control information such as the lengths of crosswalks, the locations of traffic lights, and relevant signage.

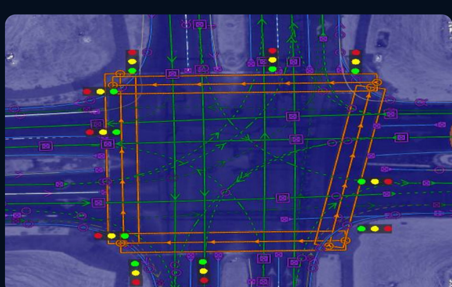

With our maps installed onboard our vehicle, our system can then focus on the parts of the environment that change dynamically around it, such as other road users. Our system can detect when a road has changed by cross-referencing the real-time sensor data with its on-board 3D map. If a change in the roadway (e.g. change in road markings) is detected, our vehicle can reroute itself within the system's operational design domain and alert our operations center to share the updated information with other vehicles in the fleet. 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. [11] 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 to a 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.

Waymo's Approach to Security

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. In addition, our software and vehicle design processes incorporate cybersecurity risk assessments, allowing us to implement defenses and protections according to the risk posed by each known vulnerability. New software releases go through an extensive peer review and verification process. Our hazard analysis, threat modeling and risk assessment processes have been designed to identify and mitigate risks that might affect safety, including those related to cybersecurity.

We use layers of security to protect our self-driving system, especially safety-critical functions like steering and braking, against unauthorized communications, including vehicle control commands. We also consider the security of our wireless communications. 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 team 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.

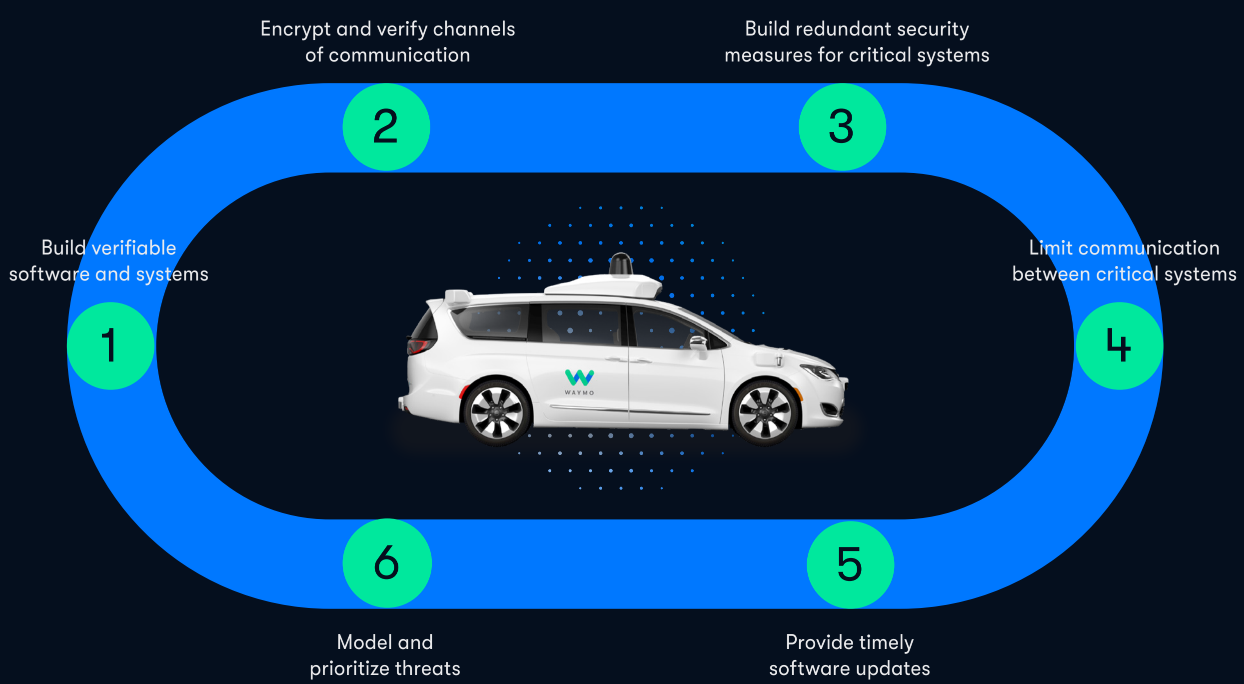

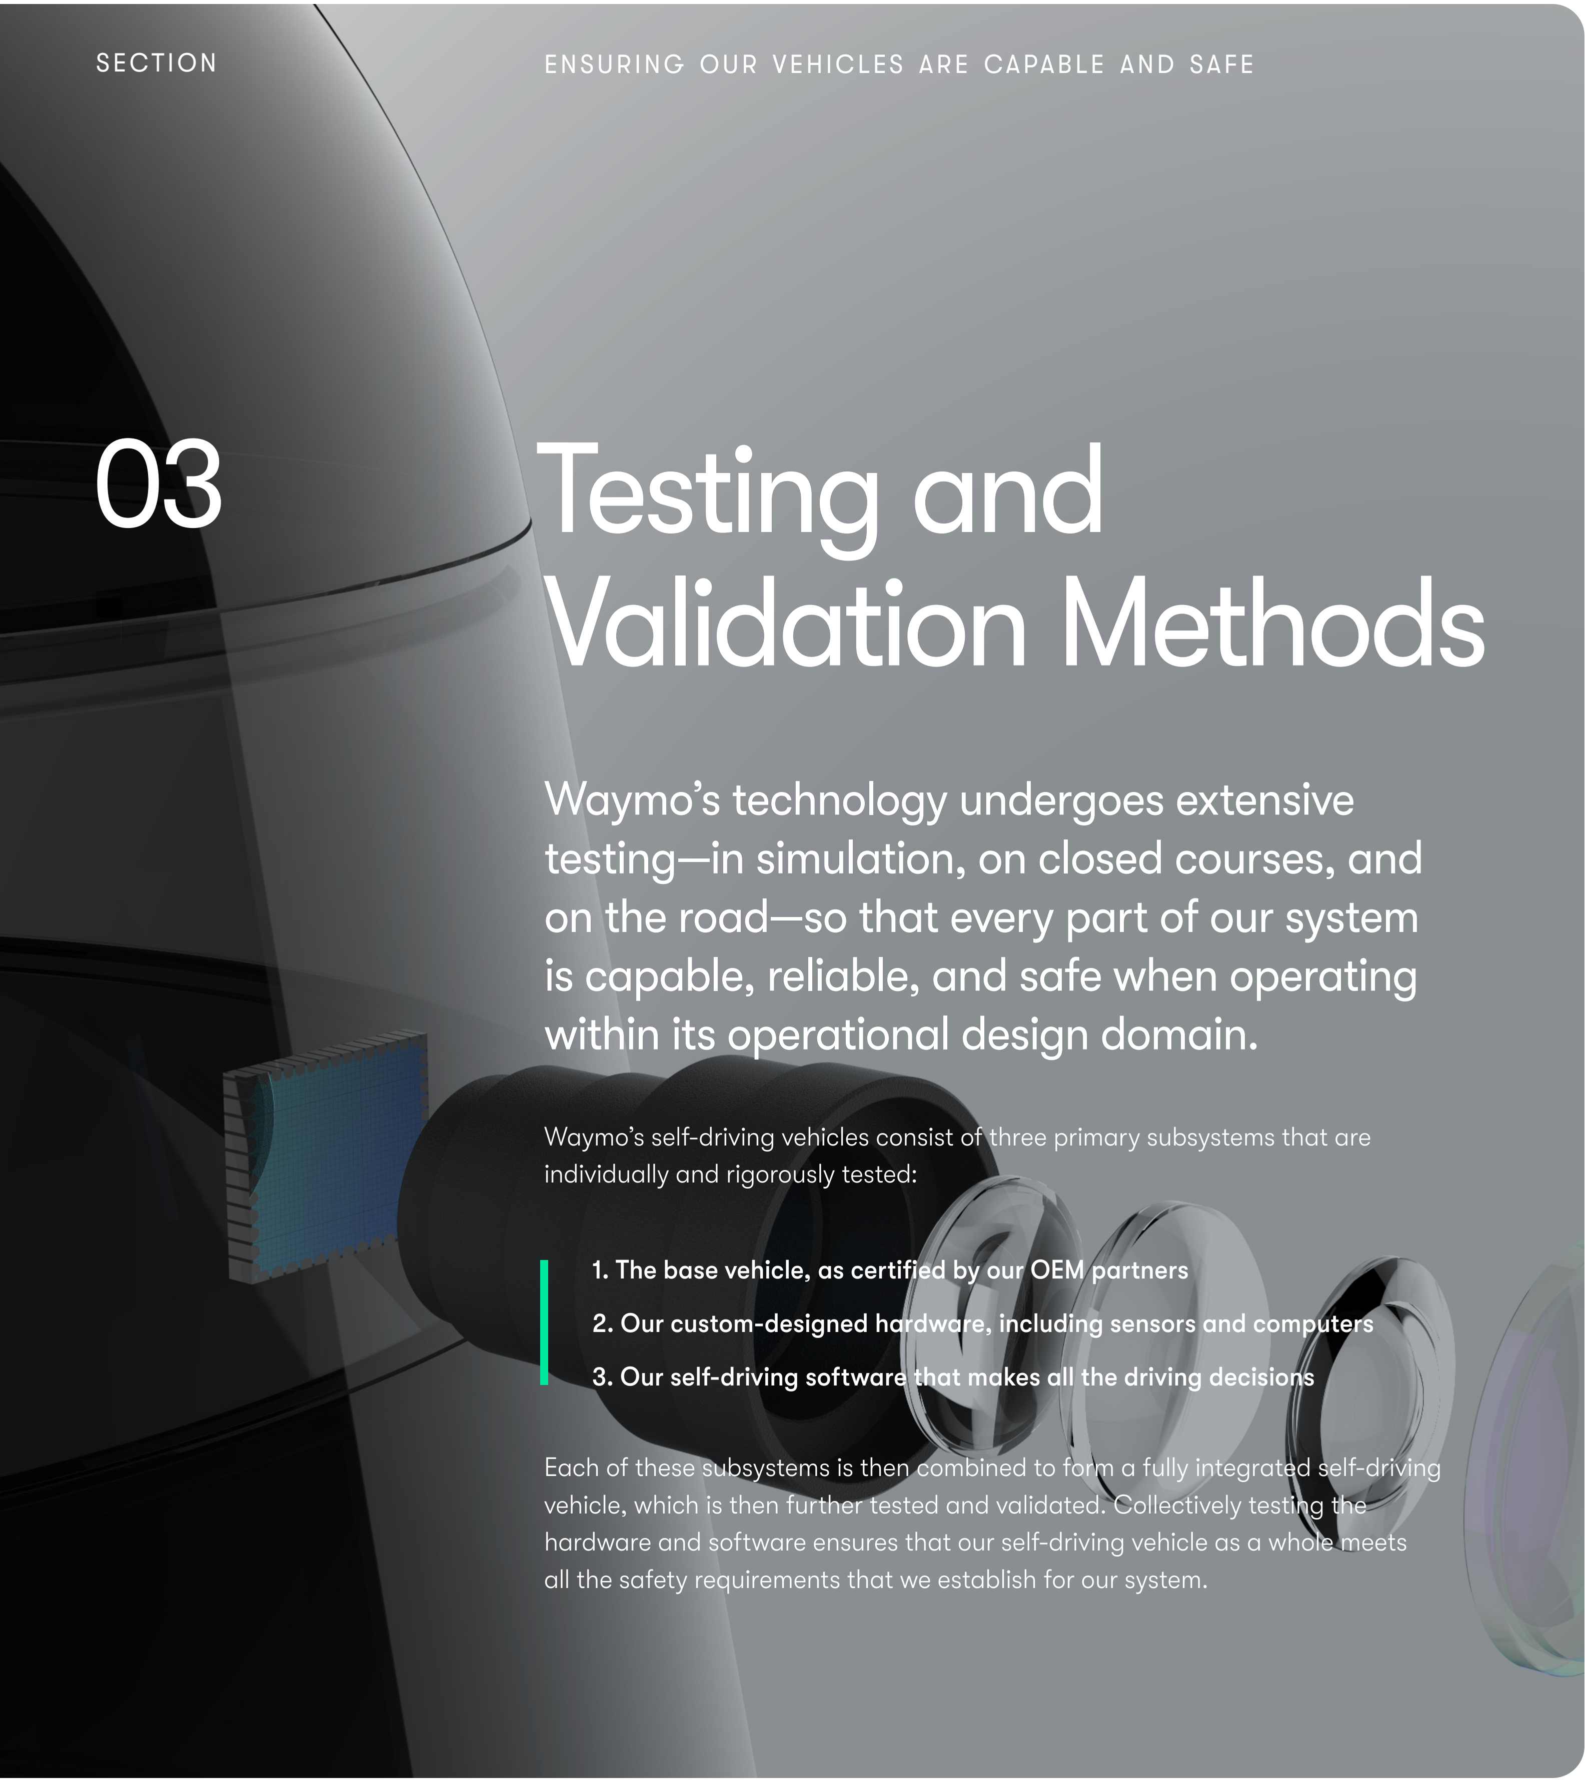

03 Testing and Validation Methods

Waymo's technology undergoes extensive testing—in simulation, on closed courses, and on the road—so that every part of our system is capable, reliable, and safe when operating within its operational 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 our OEM partners**
2. **Our custom-designed 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 self-driving vehicle as a whole meets all the safety requirements that we establish for our system.

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 evaluate suppliers' 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

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 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 trained 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.

The same Waymo Driver across several vehicle platforms

The Waymo Driver can be applied across a number of different business lines. We use the Chrysler Pacifica Hybrid Minivan for our ride-hailing service, Waymo One, and will add the Jaguar I-Pace to the fleet soon. For Waymo Via, which is focused on all forms of goods delivery, including trucking and local delivery, we use light vehicles and class 8 trucks.

Since we started testing in 2017, our Class 8 self-driving trucks have driven in a wide variety of cities and environments, from Arizona to Texas and through California and Georgia. Our trucks use the same suite of custom-built sensors as our light vehicles (configured differently), benefit from our advanced self-driving software, and leverage the over 20 million miles we've self-driven on public roads, plus the over 15 billion miles we've driven in simulation.

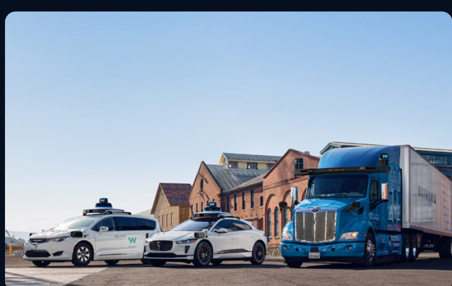

Self-Driving Software Testing

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 and 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. We also review data from crash databases and naturalistic driving studies to identify other possible collision scenarios and develop tests accordingly.

Closed-Course Testing

New software is pushed to a few vehicles first so that our most experienced drivers can test the new software, typically starting 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 and then gradually push the software update to our entire fleet after we've gained greater confidence in its performance. 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.

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

Waymo's robust simulation platform(s) enable us to replay the real-world miles we drive with each new software release, and to build completely new, realistic virtual scenarios for our software to be tested against. In simulation, we drive around 20 million miles a day, expanding the scale and complexity of our experience. This allows the Waymo Driver to refine old skills and test out new maneuvers that help it 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.

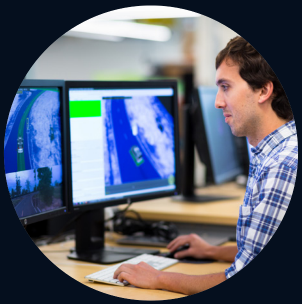

Waymo engineers build virtual scenarios that allow our self-driving vehicles to drive up to 8 million simulated miles each day.

How Simulation Works

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

01 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.

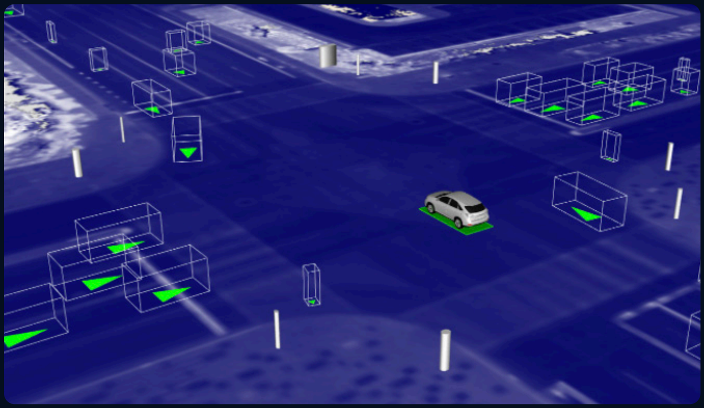

In simulation, we can practice driving the same intersection, in the same driving conditions thousands of times, with different vehicles from our fleet.

02 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.

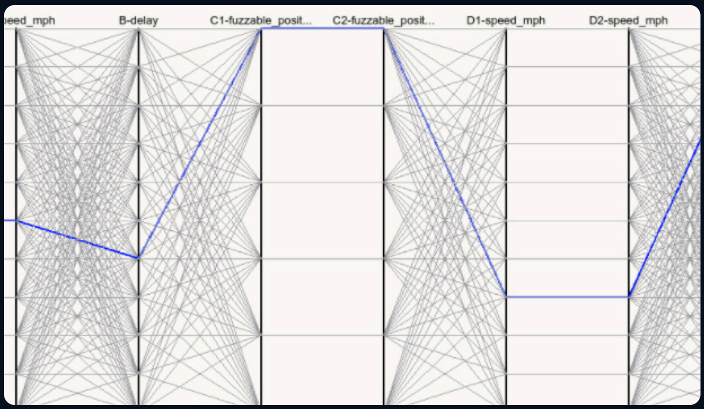

Through a process called fuzzing, we can alter the speed, trajectory and position of objects on these virtual streets.

03

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. Not only can we modify a given scenario in simulation, but we can also build entirely synthetic scenarios that we've never encountered in the real world, to understand how our system would perform.

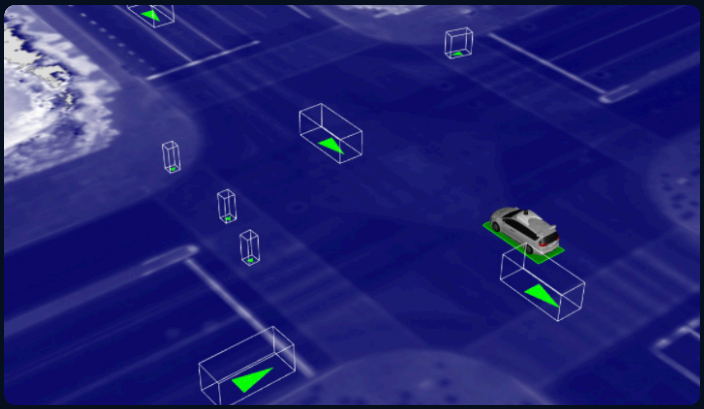

To make a scene more complex, we can add vehicles, pedestrians and cyclists that never existed in the original scene.

04

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.

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. [11] (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 at closed course facilities 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 simulation 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 vehicles, hardware, and software as an integrated fully self-driving vehicle on public roads, where it demonstrates these competencies daily in real traffic situations.

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 a decade. Driving on public roads is 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. Trained drivers undergo extensive classroom training, learning about the overall system and how to monitor the vehicle safely on public roads, practice behind the wheel on a closed course, and complete a defensive driving course. After this training, drivers are responsible for monitoring the system and if needed, taking control of the vehicle.

Our on-road testing program drives hundreds 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.

25+
Cities

10
U.S. States

Real-World Experience

Over the past decade, Waymo has tested our vehicles in ten U.S. states and self-driven in more than 25 cities—from sunny Phoenix, Arizona to rainy Kirkland, Washington, across the snowy Upper Peninsula of Michigan, through Death Valley heat, and in foggy San Francisco to ensure that our vehicles learn to drive in a variety of challenging weather conditions. As we expand to new locations, we add to our experience with different road environments, streetscapes, and human 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 drive through busy intersections and parking lots, some of the toughest pickup and drop off areas, and how to execute complex maneuvers on high-speed, 45 mph roads, such as unprotected turns. Testing in San Francisco has given us experience tackling challenging and complex situations in a dense city environment, like smoothly navigating narrow roads, double parked vehicles, and construction zones.

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.

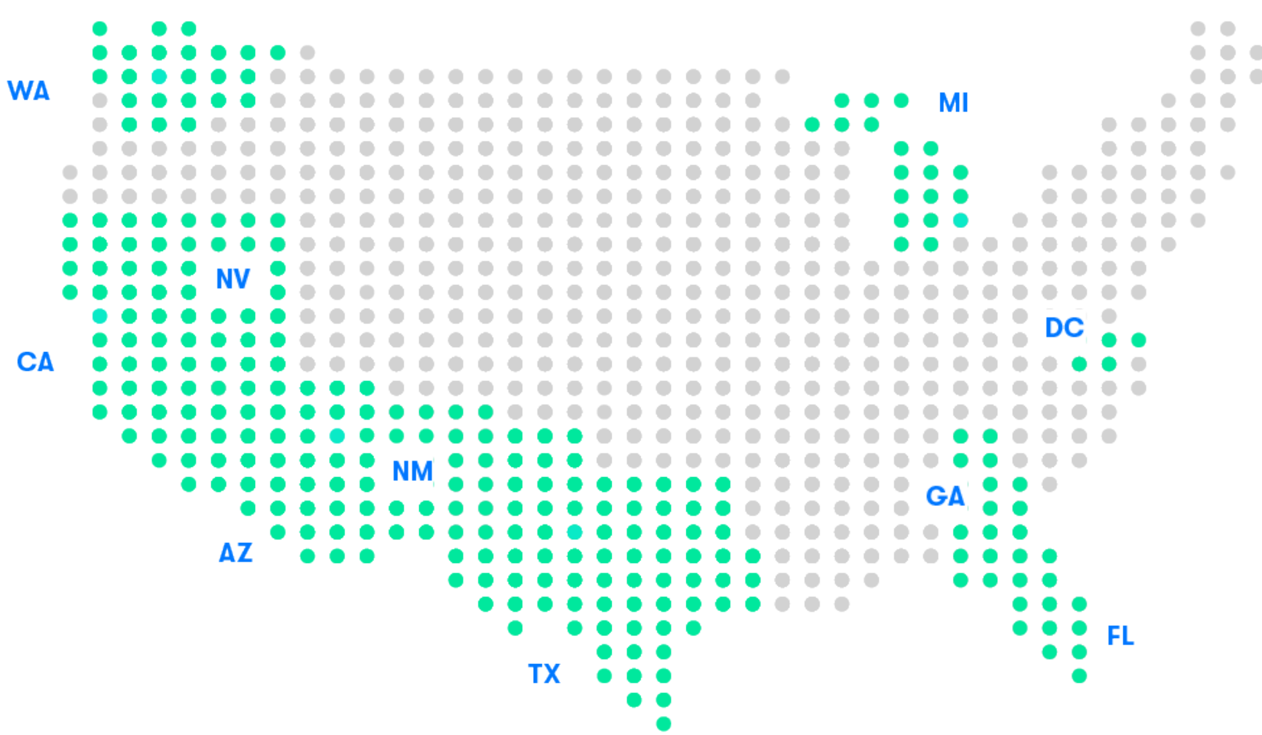

Self-Driving in Extreme Temperatures

Our goal is to enable our self-driving vehicles 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 up to and beyond the intended design capabilities.

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, using monitoring capability built into both our sensors and compute, and we constantly run software algorithms that detect anything out of the ordinary.

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 simulation 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. These sources include our own analysis of sources such as NHTSA's fatal crash database, 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. [12]

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

Field Tests at Our Closed-Course Facility

Waymo has set up a private, 113-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 simulation program, we've developed more than 40,000 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.

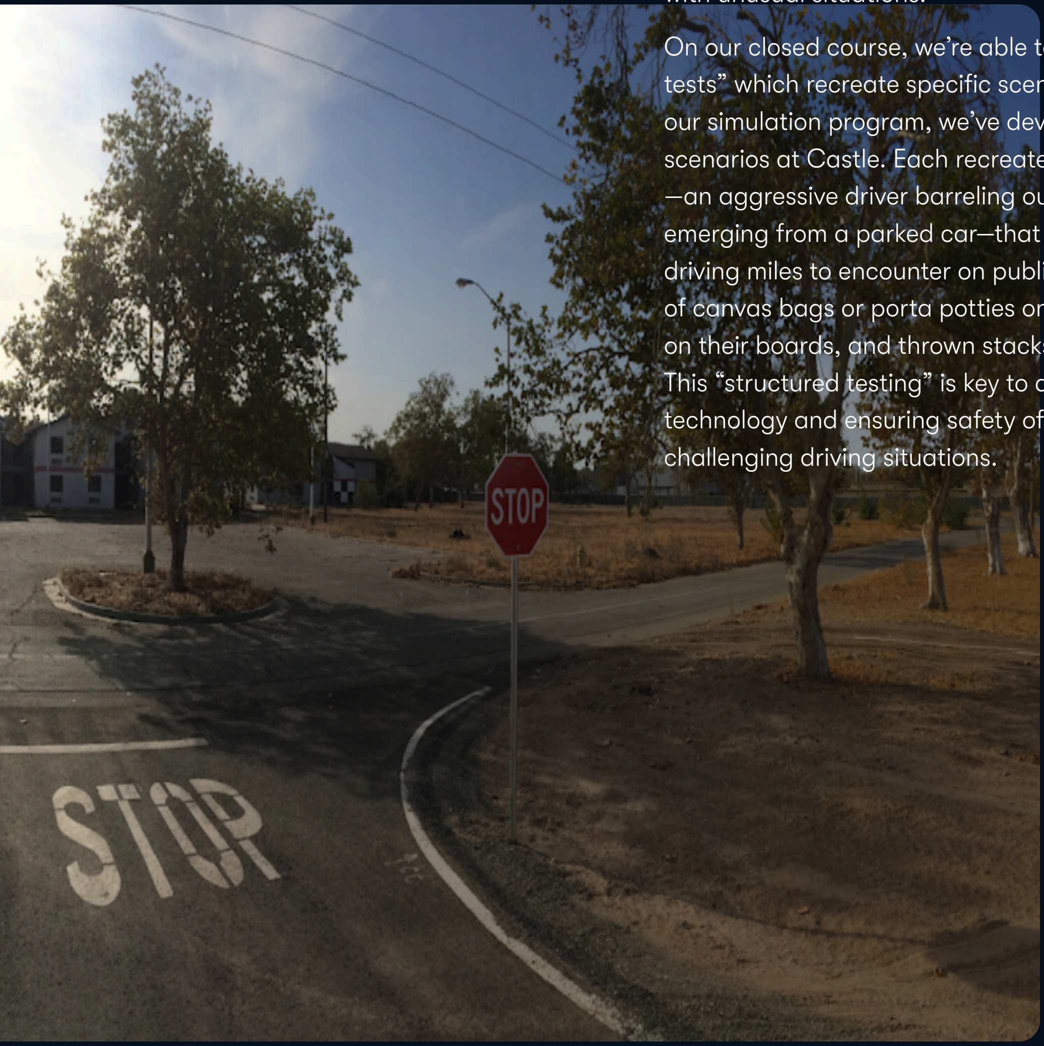

Hardware Reliability and Durability Testing

Self-driving vehicles, like their human-driven 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

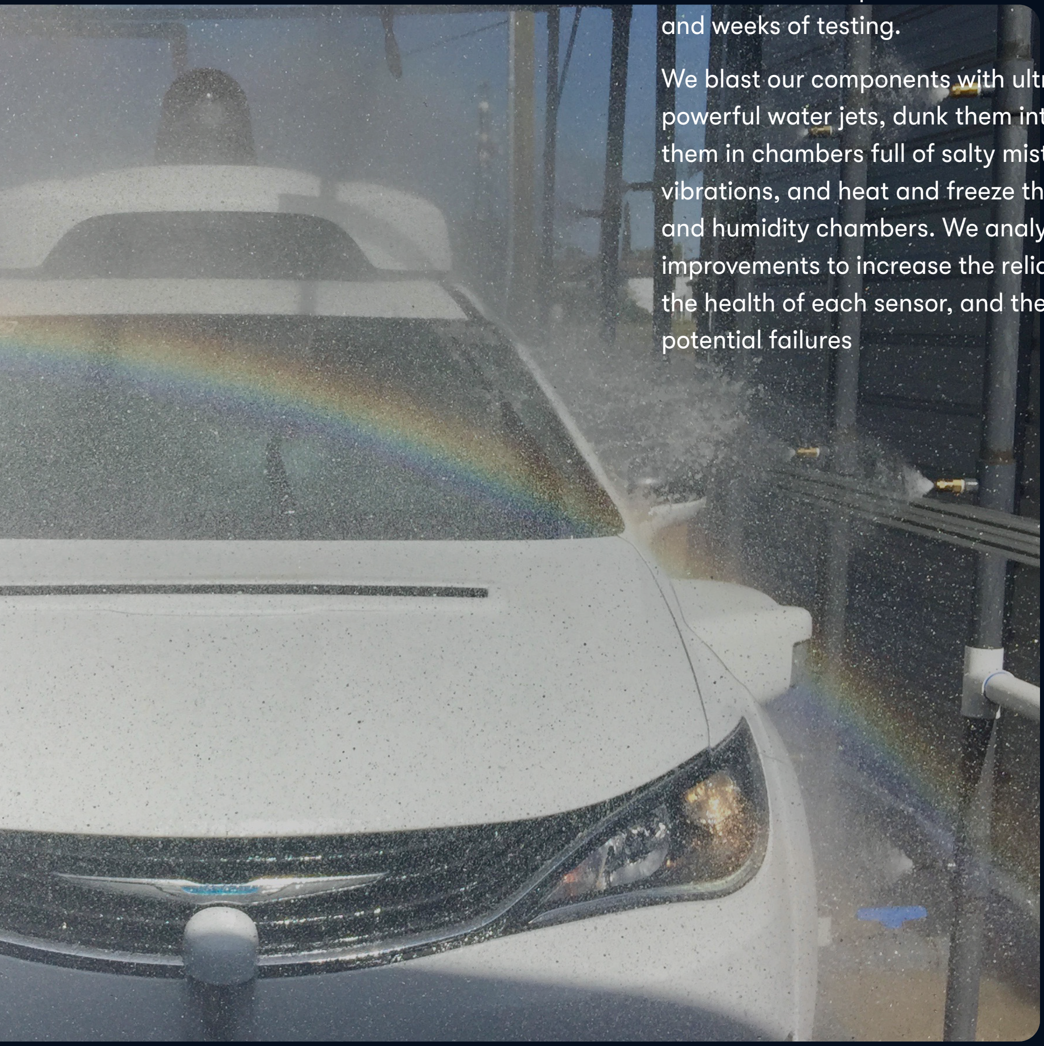

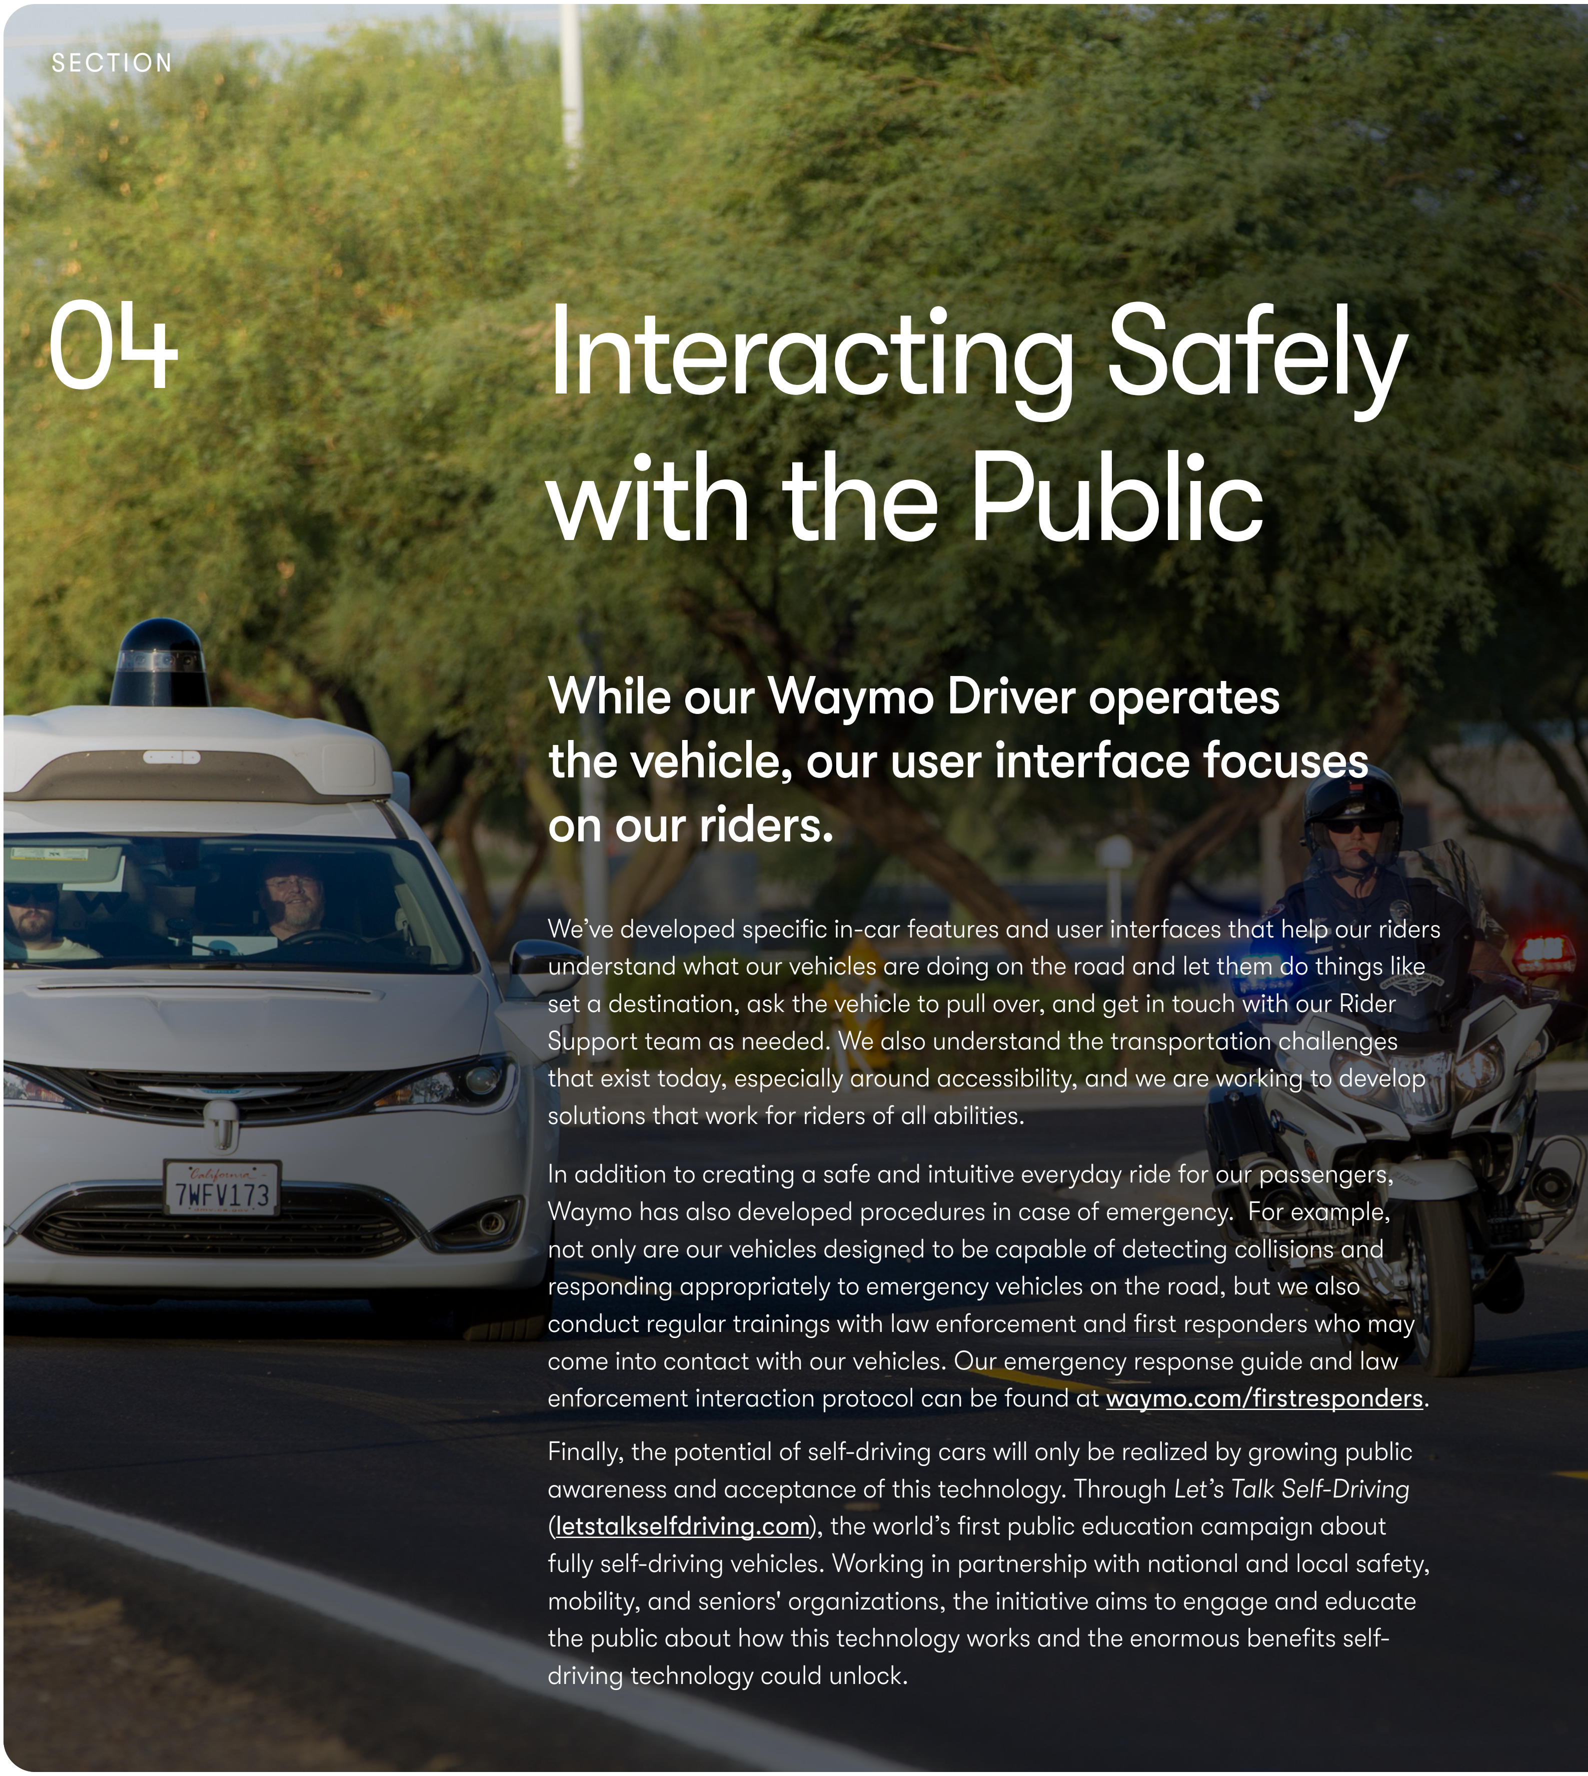

SECTION

04 Interacting Safely with the Public

While our Waymo Driver operates the vehicle, our user interface focuses on our riders.

We've developed specific in-car features and user interfaces that help our riders 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 our Rider Support team 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 be capable of detecting collisions and responding appropriately to emergency vehicles on the road, but we also conduct regular trainings with law enforcement and first responders who may come into contact with our vehicles. Our emergency response guide and law enforcement interaction protocol can be found at waymo.com/firstresponders.

Finally, the potential of self-driving cars will only be realized by growing public awareness and acceptance of this technology. Through *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' organizations, the initiative aims to engage and educate the public about how this technology works and the enormous benefits self-driving technology could unlock.

Launching the World's First Public Autonomous Ride-hailing Service

In April 2017, we first invited members of the public to ride in our cars through our early rider program. We chose a diverse and passionate group to give us feedback and work directly with our product and research teams. Their feedback helped us refine our technology and design features that meet their needs. In late 2018, Waymo became the first company in the world to launch a public self-driving ride-hailing service. We've had thousands of people using it in their everyday lives across the Metro Phoenix area. We have enjoyed seeing how they rely on the Waymo Driver to get them where they need to be, whether that's work, the grocery store, or a fun night out.

Here's what we've already learned from our riders to help improve our service.

Where They're Going:

The majority of Waymo One rides happen in the late afternoon into the evening, and our riders are using the service to do everything from commuting to work and school to running errands and visiting family and friends. Other common Waymo One use cases include going out on a date night, shopping at the local mall, or getting to the gym.

Why They Love Using It:

Not only do our riders love playing a role in shaping self-driving technology, contributing to a safer future, and the consistency of the Waymo One experience, they also enjoy the fun features we have in our car. We've added things like: dark mode to make screens dimmer at night, a zoom feature so riders can take a closer look on the screen at what the car is seeing, and a map view that allows riders to see where they are along their route. There are also functional features our riders like, including how easy it is to connect with Rider Support in the car or through the Waymo app on their smartphone.

How We Can Improve The Rider Experience:

One area we've dedicated a lot of time to improving has been pickups and dropoffs. Since there is no human driver to call, text or wave to, we have to try to make our pickup and dropoff locations as precise as possible and communicate well with riders. We've continued to make this experience smoother by making it easier to fine-tune your pickups and dropoffs in the app and improving our recommended pickup and dropoff spots.

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

Riders 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 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. Rider Support agents 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 are doing. In Waymo One, our autonomous ride-hailing service which includes our early rider program and our public service, 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 our riders.

In the event of a safety-critical event, the screens and audio system are designed to provide our riders 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.

Deploying Fully Driverless Vehicles

We've been testing in fully driverless mode since 2017 in the Metro Phoenix area—with no trained human driver behind the wheel. Since summer of 2019, we've been ramping up our driverless offerings to participants in in our early rider program in the Metro Phoenix area, and riders have been regularly matched with these fully driverless cars when they hail a Waymo vehicle using our app. By early 2020, we were serving between 1,000-2,000 Waymo One rides per week, 5-10% of which were fully driverless rides within our early rider program. We fully expect to continue to grow that number over time.

Here's how it works:

- Riders can't specifically request a fully driverless ride, but when a driverless car is nearby they may get matched with one.
  - → Our current driverless operating area in Arizona is around the same size as the city of San Francisco, which to give you an idea, is ~50 sq. miles.
- If riders do get matched to a driverless car, a notification will pop up in the riders' app letting them know.
- The car will be completely empty, no trained driver up front or anywhere in the vehicle. The space is exclusively our riders' own to kick back, relax, and do whatever they'd like without the worry of driving.
- If riders have any questions along the way, Rider Support is there to help through the tap of a button either in the app or through the in-car consoles and screens.

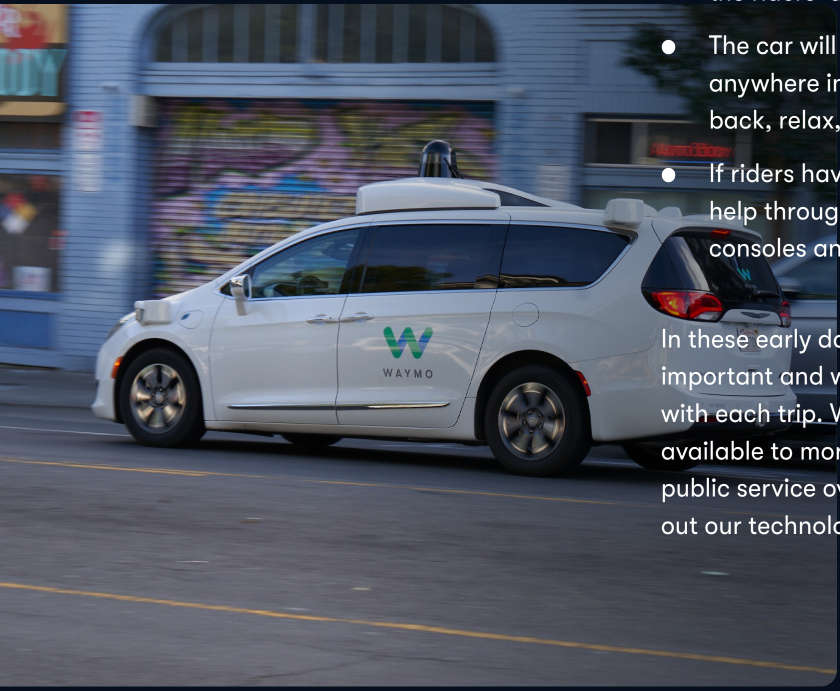

In these early days of this experience, rider feedback is especially important and we will continue to seek and use rider feedback to improve with each trip. We'll continue to learn, advance this offering, and make it available to more and more riders across both our early rider program and public service over time, just as thoughtfully and gradually as we've rolled out our technology from the very start.

Trained Drivers and Fatigue Risk Management

While fully driverless operation is our north star, trained drivers still play a critical role in our pursuit of building the World's Most Experienced Driver™. Not only do they provide vital information from the front lines to our engineering teams, they also often serve as the face of Waymo to the public and our riders. When we have a trained driver present in the driver's seat of our light and heavy duty vehicles, their most important mission is to keep themselves, our passengers, and everyone around them safe – whether they're operating for Waymo One, Waymo Via, testing on public roads, or conducting closed-course testing.

The smarter our vehicles become, the less interaction is needed from drivers. While the benefits of automation are obvious, it can actually become a problem if people get tired or bored from having too little to do. Addressing trained driver fatigue is a top priority as we work towards scaling Waymo's self-driving service. In order to manage both forms of fatigue, we have implemented a robust fatigue risk management program (FRMP), which uses state-of-the-art technologies, best practices, and input from both internal experts and leading fatigue experts at the Virginia Tech Transportation Institute, who helped review and refine our FRMP.

Elements of the FRMP include:

- First, a Psychomotor Vigilance Test evaluates each trained driver's response times to ad-hoc stimuli, to test their capacity for sustained vigilance and ability to respond effectively. Performing these tests when they arrive at work helps to ensure that their reactions are fast and accurate, which indicates that they're well rested and reduces the likelihood of fatigue during the day.
- Once inside our vehicles, we monitor trained drivers' fatigue levels.
    - → Our vehicles are also equipped with automated fatigue detection and alerts that trigger external stimuli when the system detects that a trained driver has closed their eyes or taken their eyes off the road.
- In-Car Tasks (ICT) break up the monotony of prolonged supervision and help keep drivers alert. The Waymo ICTs both prevent prolonged periods of inactivity and provide a behavioral method of assessing alertness.
- We also implemented follow-up audible alerts when a vehicle is shifted into manual mode to ensure that disengagements are intentional.
- Waymo's FRMP supplements other trained driver safety measures and applicable state and federal requirements (e.g. for heavy-duty commercial motor vehicles, hours of service, controlled substance and alcohol testing, and screening for medical conditions and sleep disorders). In addition, driver shift schedules provide multiple breaks for proper rest. Additional discretionary breaks are also suggested based on periodic fatigue surveys, which drivers take at the end of each break.

Accessibility Features In Our Self-Driving Vehicles

An Accessible Mobile App:

We've built 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 offer "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 are available in the vehicle to keep the rider informed of their journey.

Braille Labels:

The ride buttons in our self-driving Chrysler Pacificas 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 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 through our accessible app and 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 our learnings help 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 accessibility partners, including Foundation for Blind Children and the Epilepsy Foundation of Arizona, to identify solutions that can serve a broader set of individuals.

Testing with Police and Fire Departments

We've collaborated with police and fire departments in Arizona and California to conduct emergency vehicle testing with our self-driving cars. Our powerful suite of sensors, including our long-range audio detection system, observed local police vehicles, motorcycles, ambulances, fire trucks, 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 works closely with public safety officials to ensure the safe introduction of our technology in every city in which we drive. Waymo has also conducted on-site training in several cities to help police and other emergency workers identify and access our vehicle in emergency situations. Additionally, we provide detailed instructional guides, videos and lines of communication for further engagement, available at waymo.com/firstresponders.

SECTION

05

Conclusion

For more than a decade, Waymo has worked to bring 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 powered by the Waymo Driver. We are excited about the potential autonomous technology holds to improve road safety and provide new mobility options for the world. For further information about the Waymo Driver, please visit www.waymo.com.

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.

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

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 Minimal 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

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. [12]

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

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 database, and from our own extensive experience operating self-driving vehicles. [12]

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.

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.

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.

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 June 2018.

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. 2) A major part of a system which in itself has the characteristics of a system, usually consisting of several components. (MIL-STD-882E).

System. The organization of hardware, software, material, facilities, personnel, data, and services needed to perform a designated function within a stated environment with specified results. (MIL-STD-882E)

System Safety. The application of engineering and management principles, criteria, and techniques to achieve acceptable risk, within the constraints of operational effectiveness and suitability, time, and cost, throughout all phases of a system life cycle. (MIL-STD-882E).

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 the hazards or reduce the associated risks when hazards cannot be eliminated. (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.

[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 2018.” World Health Organization, 2018. https://www.who.int/violence_injury_prevention/road_safety_status/2018/en/

[3] “2018 Fatal Motor Vehicle Crashes: Overview.” National Highway Traffic Safety Administration, October 2019. https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812826; “National Highway Traffic Safety Administration's Traffic Safety Facts Annual Report” National Highway Traffic Safety Administration https://cdan.nhtsa.gov/SASStoredProcess/guest

[4] “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.

[5] Schrank, D., Eisele, B., Lomax, T., and Bak, J. “2019 Urban Mobility Scorecard.” The Texas A&M Transportation Institute and INRIX, August 2019. https://static.tti.tamu.edu/tti.tamu.edu/documents/mobility-report-2019.pdf

[6] "Fast Facts of Common Eye Disorders.” Centers for Disease Control and Prevention, June 2020. https://www.cdc.gov/visionhealth/basics/ced/fastfacts.htm

[7] 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

[8] “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

[9] “SAE J3016. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles." SAE International, June 2018.

[10] 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.)

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

[12] “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