



2018 →

Uber Advanced Technologies Group A Principled Approach To Safety

00 →

Table of Contents

01
Letter to the Reader

02
Executive Summary

03
The Future of Mobility

04
Uber's Approach to Self-Driving

05
Uber's Self-Driving Technology

06
Uber's Self-Driving Safety Principles

06.01
Proficient

06.02
Fail-Safe

06.03
Continuously Improving

06.04
Resilient

06.05
Trustworthy

01 →

Letter to the Reader



Dear Reader

At Uber, we believe that technology has the power to ignite opportunity by setting the world in motion. This is why we introduced the original Uber app in 2010. Today, you can get a ride using the Uber app in more than 600 cities across 65 countries on six continents and, in some places, you can use the Uber app to get a bike, a scooter, and connect seamlessly to public transit. It was this same confidence in the potential of technology that led us to establish our Advanced Technologies Group in 2015. We believe that introducing self-driving vehicles to the Uber digital network could make transportation safer, more efficient, and more affordable for people around the world. We believe that our efforts to develop self-driving technology are consistent with our core value to *Stand for Safety*.

In the three years since we embarked on our self-driving journey, our experiences have taught us a few valuable lessons.

First, we know that the transition toward this technology will take time. It will take time because we are committed to creating high-performance technology through rigorous software and hardware development processes. It will take time because we are committed to developing this technology with input from the people who will benefit from its availability and with governments, nonprofits, and industry groups. This means not only gathering feedback, hearing concerns, and answering questions,

but also sharing information on our progress, and seeking guidance from government stakeholders and other experts. For Uber, this is not a sprint: self-driving and human-driven vehicles will coexist on roadways for decades to come.

Next, we know that we can all benefit from this technology sooner by leveraging the depth and breadth of the Uber network and the experience that comes with running it. More than 3 million driver and delivery partners on our network enable approximately 15 million trips every day. These partners are the lifeblood of the Uber network, and they aren't going anywhere. In the early days, self-driving technology will only be able to serve some trips in some markets. As we progress in our development and look to begin connecting riders with self-driving vehicles, we will only do so when it makes the most sense for that trip. Adding self-driving vehicles to our platform could increase the size and efficiency of the Uber network as a whole, rather than replacing trips.

Third, we know that this transition is not achievable without testing on public roads. We are committed to anticipating and managing risks that may come with this type of testing, but we cannot - as no self-driving developer can - anticipate and eliminate every one.

We are deeply regretful for the crash in Tempe, Arizona, this March. In the hours following, we grounded our self-driving fleets in every city they were

Dear Reader

(continued)

operating. In the months since, we have undertaken a top-to-bottom review of ATG's safety approaches, system development, and culture. We continue to support the National Transportation Safety Board's investigation into the Tempe crash. We have taken a measured, phased approach to returning to on-road testing, starting first with manual driving in Pittsburgh. We committed to deliver this safety report before returning to on-road testing in self-driving mode, and will go back on the road only when we've implemented improved processes.

Last and most important, we know that open, regular communication with you, the public, and with other stakeholders is absolutely essential to earn your trust. Voluntary Safety Self-Assessments like this report, developed in line with the National Highway Traffic Safety Administration's guidance, will be important for facilitating public transparency and consumer education. The competitive pressure to build and market self-driving technology may lead developers to stay silent on remaining development challenges. At Uber, we believe there is extraordinary value in sharing operational safety approaches and coordinating with others in the industry to develop methods to measure and demonstrate the progress in self-driving development.

This report, and the principled approach to safety it describes, is an important step towards the greater transparency and partnership that we believe are

foundational to the success of this technology. We hope that it encourages a culture of transparency, rooted in safety, for the betterment of the industry as a whole.

Sincerely,



Dara Khosrowshahi
Chief Executive Officer

02 →

Executive Summary



Executive Summary

02 →

At a Glance

This report describes Uber Advanced Technologies Group's (Uber ATG's) approach to the safe development of self-driving vehicle technology.

This report is intended to speak to a number of audiences, including:

- **The public, fellow road-users, and potential users of self-driving technology**, who may be interested in how the technology works today, how it may fit into Uber's network over time, and how Uber is working to promote safety in its development.
- **Policymakers, including legislators, regulators, and local officials**, who may be interested in understanding the current state of the technology and the approach that Uber is taking to safety.
- **Other developers**, who may be interested in understanding the more technical aspects of our approach to safety and identifying opportunities to share information and ideas.

The technologies, policies, and programs described in this report largely reflect current capabilities of our self-driving system. Some planned capabilities are included in sections entitled *Looking Forward* in sections [06.01](#) through [06.05](#).

In *The Future of Mobility*, we put Uber's work on self-driving vehicles in the context of Uber's broader efforts to provide on-demand, multimodal transportation solutions. We describe the opportunities presented by a shift to shared, sustainable, and automated transportation, and the potential safety, mobility, economic, and environmental benefits which self-driving vehicles could provide. Achieving this transition safely will take time - to develop safe technology through safe development, earn public trust and confidence, and implement the enabling policy frameworks that will encourage best outcomes.

In *Uber's Approach to Self-Driving*, we present our mission to bring this technology to market in managed fleets of shared vehicles. We have key capabilities that support this model, including our technology and experience in ridesharing and our vibrant, established network of driver-partners. We are forming partnerships to make vehicles equipped with our own self-driving technology and vehicles equipped with other developers' self-driving technology available via the Uber platform. Underpinning these models is a fundamental belief that developing our own self-driving technology can make us more effective in safely deploying our own and other developers' self-driving vehicles.

In *Uber's Self-Driving Technology*, we provide an overview of the base vehicle that is the foundation for our self-driving system, the hardware elements which we integrate into the base vehicle to deliver data inputs to the self-driving computer, and the software processes which take these data inputs to perceive, predict, plan, and execute the vehicle's movement.

In *Uber's Self-Driving Safety Principles*, we introduce our safety case approach. The U.S. Department of Transportation and its National Highway Traffic Safety Administration (NHTSA) identify 12 safety elements, or core areas of consideration with respect to safety and self-driving. Our approach is to stitch these together into a safety case: a convincing and comprehensive argument that our self-driving system is appropriately safe to operate.

A successful safety case convinces stakeholders that the risk of harm from a system has been reduced to an acceptable level. It does this by analyzing risk for a given Operational Design Domain (ODD) and establishing an overall premise for system safety in a set of principles, expressed as requirements. These principles are then used to inform the development

Executive Summary

02 →

At a Glance (continued)

of more detailed requirements, which are allocated to development processes and substantiated.

The remainder of the report describes Uber’s five self-driving Safety Principles. We believe that for a self-driving vehicle to be acceptably safe to operate, it must be shown to be:

- **Proficient** - In the absence of hardware faults, how do we demonstrate that our system performs more safely than human drivers, using credible and tractable performance metrics?
- **Fail-Safe** - How do we ensure that the system responds to a malfunction that could result in harm to a person by transitioning to a state which reduces the risk of harm? Under what circumstances will we allow a risk to persist?

- **Continuously Improving** - How do our development processes capture, consider, and respond to undesirable or unexpected system behavior?
- **Resilient** - How do we prevent, protect, and/or warn against potential harm that arises when our technology is used counter to its design or purpose by external actors?
- **Trustworthy** - How do we create and maintain a two-way dialogue with our riders, regulators, legislators, other road users, and advocacy organizations and provide them with evidence of safe performance?

Each principle addresses a number of NHTSA’s safety elements, as summarized in the diagram below. □

NHTSA Safety Elements and Safety Principles Crosswalk

Key

- P1 – Principle 1: Proficient
- P2 – Principle 2: Fail-Safe
- P3 – Principle 3: Continuously Improving
- P4 – Principle 4: Resilient
- P5 – Principle 5: Trustworthy

	System Safety	Operational Design Domain	Object & Event Detection & Response	Fallback (Min Risk Condition)	Validation Methods	Human Machine Interface	Vehicle Cyber-security	Crash-worthiness	Post-Crash ADS Behavior	Data Recording	Consumer Education & Training	Federal, State, & Local Laws
P1	✓	✓	✓									
P2	✓			✓								
P3					✓							
P4						✓	✓	✓	✓	✓		
P5	✓										✓	✓

03 →

The Future of Mobility



The way that people and goods are transported today does not need to be the way they are transported in the future.

Globally, an enormous number of cars are produced every year.

— According to the International Organization of Motor Vehicle Manufacturers (OICA), there were nearly one billion passenger cars in use globally in 2015¹ and an additional 70 million were sold in 2017.²

These cars are rarely used.

— A 2012 report by the UK-based RAC Foundation³ reaffirms parking guru Donald Shoup's earlier assertion that passenger cars remain parked more than 95 percent of the time.⁴ Since these cars are rarely used, they take up a lot of space in our cities when they sit parked for the majority of the day. A 2010 study by researchers at the University of California,

Berkeley assumes that there are between 3.4 and eight parking spaces per passenger vehicle in the U.S., taking into account differences between urban and rural settings.⁵

These cars are used inefficiently.

— In 2017, 60 percent of passenger car miles travelled in the U.S. were driven by the driver alone.⁶ Since these cars are used inefficiently, city streets are congested at peak commuting hours with large vehicles and often single occupants. This leads to lost time, increased stress, and reduced productivity. Energy consumption and emissions per passenger-mile are high when vehicles are utilized in this way.

At Uber, we believe that the future of mobility is increasingly shared, automated, and sustainable. □

¹ OICA, ND, 'Passenger Cars World Vehicles in Use.'

² OICA, ND, 'Provisional New Passenger Car Registrations or Sales.'

³ RAC Foundation, 2012, 'Spaced Out: Perspectives on Parking Policy.'

⁴ Shoup, 2005, 'The High Cost of Free Parking.'

⁵ Chester, et al., 2010, 'Parking infrastructure: energy, emissions, and automobile life-cycle environmental accounting.'

⁶ Uber analysis of U.S. DOT Federal Highway Administration, 2017, 'National Household Travel Survey.'

Shared

Sharing vehicles at scale can make transportation:

—

Less expensive relative to personal car ownership by:

- Centralizing and sharing insurance, maintenance, and parking costs.
- Distributing the cost of a single ride between riders.

More convenient relative to personal car ownership by avoiding the time and cost associated with parking.

In four out of the five largest cities in the U.S., it may already be more cost-effective to share a car than to buy, maintain, and park a personal vehicle.⁷ Despite this, today, less than 1 percent of passenger miles traveled are carried out using shared car services.^{8,9}

Uber enables sharing:

—

Of a single car by multiple users over the course of a day through UberX

If the same car can be shared to satisfy the needs of multiple users throughout the day, the number of total cars needed per person goes down.

Of a single car by multiple users at the same time through Uber POOL

Separately, if we can make it appealing for commuters to share the same vehicle at the same time through increased convenience and reduced cost, fewer vehicles will need to share the same roadway at the same time. This could reduce congestion on our roads and energy consumption and emissions per passenger mile.

Companies like Uber will continue to develop their ridesharing networks, improve coverage and reliability, and find new ways to encourage pooling. □

⁷ Kleiner Perkins, 2018, '[Internet Trends Report](#).'

⁸ McKinsey & Co., 2017, '[The Automotive Revolution is Speeding Up: Perspectives on the Emerging Personal Mobility Landscape](#).'

⁹ Uber, 2018, '[Three Early Takeaways from the 2017 National Household Travel Survey](#).'

Sustainable



Motor vehicle transportation creates local air pollution and contributes to climate change. In 2017, carbon emissions from energy consumption in transportation were higher than any other sector in the U.S. – emissions from transportation grew in each of the last five years while emissions fell in every other sector.¹⁰ However, promising trends in fuel efficiency and low- and zero-emission vehicle adoption have the potential to reverse this trend:

- According to the U.S. Environmental Protection Agency (EPA), in 2016, average new vehicle carbon-dioxide emissions per mile reached a record low, and fuel economy reached a record high.¹¹ These improvements in fuel efficiency and fuel economy in new vehicles mean that fleets comprised of newer vehicles will reduce environmental impact relative to older, less efficient ones.

- According to the International Energy Agency (IEA), the number of battery electric vehicles (BEVs) and plug-in hybrid electric vehicles (PHEVs) on the world's roads exceeded 3 million in 2017, a 54 percent increase compared with 2016.¹² IEA attributes this growth to government policies and continued improvements in the performance and cost of battery technologies. The decreased cost and increased reliability of BEVs and PHEVs have the potential to further reduce the environmental impact of transportation.

These technologies may be most readily incorporated into fleets of shared vehicles because higher efficiency vehicles have lower operating costs and vehicles utilized more intensively are replaced more quickly. □

¹⁰ U.S. Energy Information Administration, 2018, 'Monthly Energy Review: Environment,' August 2018 report, retrieved 2 September 2018.

¹¹ U.S. EPA, 2018, 'Light-Duty Automotive Technology, Carbon Dioxide Emissions, and Fuel Economy Trends: 1975 Through 2017.'

¹² IEA, 2018, 'Strong policy and falling battery costs drive another record year for electric cars.'

Automated

Many studies have sought to estimate the impacts of self-driving technology on the economy, the environment, urban design and parking requirements, employment, and road safety outcomes. Taken together, these studies describe what may feel like an unimaginable future world where most or all cars are driverless and most or all rides are shared.

At Uber, while we are excited by these transformational changes and the future state, we are focused on the ways that self-driving technology can create incremental positive changes in everyday life for the communities we serve.

We believe that automated driving technology can be:

— Safer

Self-driving vehicles hold the potential to drive more safely than a human driver. Computers can look in all directions at once, and they don't get distracted, fatigued, or impaired.

More cost-efficient

Operated in shared fleets at scale, self-driving cars can be cheaper to operate than human-driven cars, improving the economics of ridesharing relative to personal car ownership.

More time-efficient

Riders who now spend time driving on congested freeways can reclaim this time for work or leisure. If sharing reduces congestion, these riders can also have shorter commutes.

More space-efficient

As more people share rides, the number of parking spaces required could fall, parking lots could shift out of cities to make room for other uses, and curb space may need to be more efficiently allocated.

More equitable than existing transportation options

Shared, automated mobility can work to extend the reach of public transit and bridge the first/last mile gap in areas typically underserved by transit systems, and for certain populations like people with disabilities, youth, and seniors.

Better for the environment

When combined with automated driving technology, appropriate policies that incentivize sharing, improve fuel efficiency, and discourage driving without any passengers have the potential to take cars off the road.^{13, 14, 15} □

¹³ Greenblatt & Saxena, 2015, 'Autonomous taxis could greatly reduce greenhouse-gas emissions of US light-duty vehicles.'

¹⁴ International Transport Forum, 2016, 'App-Based Ride and Taxi Services: Principles for Regulation.'

¹⁵ University of California, Davis & Institute for Transportation and Development Policy, 2017, 'Three Revolutions in Urban Transportation.'

Getting There

Like other self-driving technology developers, we are excited for this future, and we are working to make this technology and its potential advantages a reality that benefits all road users. But, we know that it will - and it should - take time to transition toward a future characterized by shared, sustainable, automated transportation.

In particular, it will take time to:

– **Develop Safe, Reliable Self-Driving Technology**

Developers are making incredible progress towards automation for small, constrained Operational Design Domains (ODDs); significant development work remains to deliver on automation for a wider range of ODDs.

Safely Develop Self-Driving Technology

In addition to designing for safe performance, we are committed to undertaking our development efforts carefully and keeping our testing team safety-focused. We do this by developing and enforcing safety-forward training and operational procedures.

Earn Public Confidence

Competitive pressures have made sharing information on progress in development challenging. Yet transparency into developments and progress are important to earn and increase public confidence in this technology and, in turn, its ability to deliver on the potential benefits.

Design and Implement Supportive Policies

Policymakers are focused on providing appropriate frameworks for development and testing; over time, focus will shift toward policies which support the transition toward this shared, sustainable, and automated future, including, e.g. infrastructure investment, appropriate road use pricing, retraining and workforce support programs. □

04 →

Uber's Approach to Self-Driving



Uber's Approach to Self-Driving

Our Mission

We believe the best way to harness the power of self-driving technology for broad public benefit is to deploy it in managed fleets of shared vehicles equipped with Level 4 capability.¹⁶

Delivering self-driving technology through this model makes it possible for fleet operators to use vehicles more efficiently and realize economies of scale associated with fleet-based maintenance and repair. This fleet-based approach can also manage risks associated with personal ownership of self-driving vehicles, especially the risk that reduced cost of travel leads to more travel and more congestion. □

Our Key Capabilities

Uber is positioned to successfully develop and deploy self-driving technology because:

— **Our technology and experience in ridesharing have prepared us to own and operate a world-class fleet of self-driving vehicles or to match up third-party fleets with riders.**

We already have established systems which evaluate a rider's needs and connect them with a driver-partner who can best meet these needs. Driver-partners operating on our network have driven most roadways in the cities where we operate, and we have learned to anticipate and plan for areas and times of high demand. We understand how to develop human-centered technology products to meet personal transportation needs. All of this direct experience and information positions us to be a next-generation fleet operator, with the expertise to apply a wealth of data to fleet operation needs, such as maintenance, monitoring, customer service, and rider experience, as well as a platform for matching up third-party fleets with Uber riders.

Our vibrant, established network of driver-partners can work together with self-driving vehicles to deliver a consistent, reliable rider experience.

We expect our self-driving vehicles to complement our existing products. Initially, self-driving vehicles will only safely serve some trips under some conditions, based on road geometry, weather, and other factors. Because we intend to integrate self-driving vehicles into our wider network, we will not need to send a self-driving vehicle to places where it is not yet prepared to go. □

¹⁶ A vehicle equipped with Level 4 technology is capable of "high driving automation," or "sustained and ODD-specific performance by an [Automated Driving System] of the entire [Dynamic Driving Task (DDT)] and DDT fallback without any expectation that a user will respond to a request to intervene." See SAE International, 2018, 'J3016_201806: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles.'

Uber's Approach to Self-Driving

Our Partnerships

There are a number of different pathways by which shared, self-driven fleets can be safely and efficiently brought to market.

Our self-driving strategy centers around partnership, because we know that extremely valuable experience in automotive manufacturing abounds. By combining Uber's self-driving technology with partners' state-of-the-art vehicles and production capacities, we'll get to the future faster than going it alone.

We are actively developing partnerships in two models:

—

- 1. Fleets of vehicles equipped with our own self-driving technology made available via the Uber platform.**

Under this model, we develop and validate our own self-driving technology, both hardware and software, and work in collaboration with an Original Equipment Manufacturer (OEM) to integrate our technology into a base vehicle. We have been working in partnership with Volvo to pilot this model by integrating our self-driving technology into Volvo XC90s.

- 2. Fleets of vehicles equipped with other developers' self-driving technology made available via the Uber platform.**

Partners develop and validate their own self-driving technology, either as a purpose-built vehicle or integrated into a base vehicle. These vehicles may be owned and operated by the partner or a third-party fleet manager. These vehicles will need to meet a number of criteria related to safety

and user experience before being hosted on the Uber network; our approach to developing these standards is covered in [section 06.05](#).

Underpinning these models is a fundamental belief that developing our own self-driving technology can make us more effective in safely deploying our own and other developers' self-driving vehicles. □

05 →

Uber's Self-Driving Technology



Uber's Self-Driving Technology

05 →

Base Vehicle

Uber selects vehicle platforms with a strong track record of safety and high marks in passive safety testing by independent ratings agencies.

All of the vehicles in Uber's current fleet are recent model-year Volvo XC90 sport-utility vehicles, upfitted with sensors and our self-driving technology. The XC90 has been recognized as one of the safest vehicles in the world.^{17,18} The 2017 and 2019 models were Insurance Institute for Highway Safety's (IIHS's) Top Safety Picks.¹⁹

Key Safety Features of the Volvo XC90²⁰

— These features are available on both our current and future generation vehicles.

City Safety Automatic Emergency Braking (AEB) System

The AEB is a diverse sensing and compute software system which operates independently of Uber's self-driving system. The AEB includes a forward facing radar, camera, and Electronic Control Unit (ECU).

Anti-Lock Braking System (ABS)

The ABS helps to improve vehicle control during braking by automatically modulating to help prevent lockup.

Seatbelts with Pre-tensioners and Load Limiters

Pre-tensioners tighten safety belts in the event of a collision and load limiters minimize belt-inflicted injury.

Electronic Stability Control (ESC)

ESC consists of traction control, spin control, active yaw control, and engine drag control. It helps to reduce wheel spin, counteract skidding, and improve directional stability. □

¹⁷ Volvo Car Group, 2016, '[Volvo XC90 wins North American Truck of the Year – again.](#)'


¹⁸ Automotive World, 2018, '[Volvo XC90 is a genuine life-saver.](#)'

¹⁹ IIHS, ND, '[2017 Volvo XC90.](#)'

²⁰ Volvo Car Group, ND, '[Volvo XC90 Features.](#)'

Hardware

Current Generation



1. Light Detection and Ranging (LIDAR)

LIDAR is a remote sensing method that uses light in the form of a pulsed laser to measure distances to actors and objects. Each upfitted XC90 is equipped with one, top-mounted LIDAR unit. Uber's self-driving system utilizes a LIDAR unit with a range of over 100 meters (m).

2. Cameras

Each upfitted XC90 is equipped with cameras that provide high resolution, near-, medium-, and long-range imagery. There are cameras mounted in the sensor pod on top of the vehicle and around the vehicle for 360° coverage. The camera hardware and accompanying firmware are custom to the Uber self-driving system. Some of these cameras have a wide field of view and some have a narrow field of view.

A system of cameras provides imagery to support near-range sensing of people and objects within 5m from vehicle, in particular to assist during pick up and drop off, lane changing, and parking.

3. Radar

Each upfitted XC90 is equipped with radars that provide object detection, ranging, and relative velocity of objects. Forward-facing radars are mounted below the headlamps, side-facing radars are mounted in the front and rear corners of the vehicle, and rear-facing radars are mounted near the ends of the bumper beam.

4. Global Positioning System (GPS)

The GPS system provides rough position to support

localization, vehicle command, map data collect missions, and satellite measurements.

5. Self-Driving Computer

The self-driving computer is the main system computer running Perception, Prediction, Motion Planning, and other software. The computer hardware and firmware are custom to Uber's self-driving system. The computer is liquid-cooled for high power heat rejection.

6. Telematics

Custom telematics hardware and software provide cellular data communication to support carrier network redundancy, secure mobile data traffic, and authenticated cloud communication. □

Hardware

Next Generation



In addition to the elements described on the prior page, we intend for the next generation of our self-driving vehicles to be equipped with:

1. Ultrasonic Sensors (USSs)

USS provide near-ranging sensing of people and objects within 5m from the vehicle, in particular to assist with stopping and starting, lane changing, and parallel parking.

The USS will use echolocation to range objects. These sensors will be distributed across the front and rear fascia and the starboard and port side sills.

2. Vehicle Interface Module (VIM)

The VIM is a gateway to allow the self-driving computer to communicate with the various vehicle control systems. It has been developed in accordance with International Organization for Standardization (ISO) 26262 Automotive Safety Integrity Level D (ASIL-D)²¹ and provides closed-loop motion control, undertaking both trajectory management and trajectory tracking. The VIM is designed to be fully redundant. Its onboard inertial measurement units (IMUs) enable the VIM to safely navigate the vehicle to a stop in the event of certain autonomy system faults. □

²¹ ISO, 2011, 'ISO 26262 Functional Safety for Road Vehicles.'

Uber's Self-Driving Technology

Software

The self-driving system must be capable of detecting and responding to a variety of static and dynamic actors and objects in the road environment. Sensing hardware, including LIDAR, cameras, and radars, generate input data for the vehicle's software system. The vehicle software uses that input data to observe and categorize actors and objects in the environment, predict the actions of the actors and objects it finds, and then plan a safe path for the vehicle premised on the rules of the road.

In addition to data generated by the sensor suite described in the prior section, our self-driving software uses a set of our high-definition maps, which we develop to improve real-time understanding of the driving environment. Our maps include the following information layers, among other data:

- Geometry of the road and curbs
- Drivable surface boundaries and driveways
- Lane boundaries, including paint lines of various types
- Bike and bus lanes, parking regions, stop lines, crosswalks
- Traffic control signals, light sets, and lane and conflict associations
- Railroad crossings and trolley or railcar tracks
- Speed limits, constraint zones, restrictions, speed bumps
- Traffic control signage

Our self-driving software uses these high-definition maps together with the real-time data delivered by the onboard sensors described in the previous section to perform a series of automated tasks:

Perception

Our self-driving system features an array of overlapping sensors gathering data covering 360°

around the vehicle. Our Perception software processes this data and combines it with maps into a full representation of the environment.

Localization

Inputs from our sensor suite and our high-definition maps allow our self-driving system to determine precisely where it is in the world, down to within a few centimeters.

Prediction

Our Prediction software takes a representation of the driving environment from Perception and uses this representation in order to predict what the actors or objects in the environment are likely to do next.

Routing and Navigation

Our Routing software leverages high-definition map data, vehicle status, and operational activity to determine what route the vehicle should use to reach its intended destination. Depending on the type of mission and vehicle state, the route generated may be constrained for operational accuracy or optimized for operational efficiency.

Motion Planning

Our Motion Planning software takes into consideration the built environment and mapped information, the generated route plan, and inputs from Perception and Prediction to generate a motion plan for the vehicle.

Vehicle Control

Our Vehicle Control software executes the trajectory supplied by Motion Planning by controlling the actuation of the vehicle and driving direction through communication interfaces. □

Uber's Self-Driving Technology

05 →

The Brain of an Uber Self-Driving Vehicle



High-Definition Map and Localization


Framework

Computing



06 →

Uber's Self-Driving Safety Principles



Uber's Self-Driving Safety Principles

Our Safety Case

Uber's vision for self-driving technology is built on a foundational commitment to safety. We develop our technology through an iterative cycle, including both virtual and real-world testing. For this technology to become highly proficient and reliable, it must share the road with the public when this can be done safely and responsibly. It is our responsibility to ensure that we are developing and deploying this technology in a manner that does not introduce undue risk to the public. We must be confident our self-driving system is capable of operating safely on public roads long before it ever gets there.

In order to demonstrate this readiness, we are creating a safety case: a convincing and comprehensive argument that our self-driving system is appropriately safe to operate. A successful safety case convinces

stakeholders that the risk of harm from a system has been reduced to an acceptable level.²² To properly understand and assess risk, we first establish an ODD. An ODD defines the intended usage of the system and the likelihood and severity of hazards in that context. From this, we then establish an overall premise for system safety by identifying requirements, allocating responsibility, and providing substantiating evidence and artifacts ensuring the plan is correctly implemented.

²² Kelly, 2004, 'A Systematic Approach to Safety Case Management.'

Our Self-Driving Safety Case

Presents an analysis of risk (ODD)

- Defines the intended usage of the system, and studies the likelihood and severity of hazards in that context.

Establishes the overall premise for system safety (Safety Principles)

- Identifies the necessary and sufficient set of safety principles and demonstrates conformance on these principles.

Identifies requirements that must be fulfilled to meet this premise

- Establishes exactly what the development process and each part of the system needs to achieve to ensure safety.

Specifically allocates responsibility to individual elements

- Assigns ownership of all development processes ensuring responsibility at each step.

Details how conformity to this scheme will be substantiated

- Shows how configuration management, testing, and validation ensures the plan is correctly implemented.

Uber's Self-Driving Safety Principles

06 →

Our Safety Case (continued)

To anchor our safety case, we began developing a set of self-driving vehicle Safety Principles in 2016. This set of necessary and sufficient high-level criteria governs our safe development and deployment of this technology - to fulfill these principles requires both rigorous system development and dependable organizational processes. We believe that for a self-driving vehicle to be acceptably safe to operate, it must be shown to be:


06.01 Proficient

06.02 Fail-Safe

06.03 Continuously Improving

06.04 Resilient

06.05 Trustworthy



Uber's Self-Driving Safety Principles

Our Safety Case

NHTSA Safety Elements and Safety Principles Crosswalk

Key

- P1 – Principle 1: Proficient
- P2 – Principle 2: Fail-Safe
- P3 – Principle 3: Continuously Improving
- P4 – Principle 4: Resilient
- P5 – Principle 5: Trustworthy

	System Safety	Operational Design Domain	Object & Event Detection & Response	Fallback (Min Risk Condition)	Validation Methods	Human Machine Interface	Vehicle Cyber-security	Crash-worthiness	Post-Crash ADS Behavior	Data Recording	Consumer Education & Training	Federal, State, & Local Laws
P1	✓	✓	✓									
P2	✓			✓								
P3					✓							
P4						✓	✓	✓	✓	✓		
P5	✓										✓	✓

We intend for these principles to remain at the core of our development and operational efforts, though the approach to fulfilling each principle may change over time. For example, today we rely on our Mission Specialists to take over in certain situations that we anticipate the self-driving system will handle independently in the future.

Additionally, as other developers seek to answer similar questions, e.g. the method of measuring the safety advantage of self-driving vehicles over human drivers, we expect that formal and informal standards will emerge, including industry best practices and/or standards required by law. These standards governing self-driving vehicle safety will be reshaped and refined through a process that includes other developers and key stakeholders.

These self-driving Safety Principles ground our holistic approach to safety during development efforts and ensure safety is ingrained in each step of the process, from initial concept through vehicle end of life. Guidance from the U.S. Department of

Transportation (DOT) and its National Highway Traffic Safety Administration (NHTSA) identifies 12 safety elements, or core areas of consideration with respect to safety and self-driving.^{23, 24} Our Safety Principles encompass all of these safety elements; each element is represented within at least one, if not each, principle.

→ **Looking Forward**

While this report is primarily focused on where we are today, it is important to be mindful of where we are heading. Our approaches, design features, and procedures discussed here are focused on our current capabilities and developmental process. Doing so not only provides a clearer picture of the possible road ahead and path to scalability, but it also helps to better contextualize our development processes and methodologies. Where applicable, future-facing plans are discussed in supporting *Looking Forward* sections. □

²³ NHTSA, 2017, 'Automated Driving Systems 2.0: A Vision for Safety.'

²⁴ NHTSA, 2018, 'Preparing for the Future of Transportation: Automated Vehicles 3.0.'

Principle 1

Proficient



In order to improve safety, self-driving systems must, at a minimum, perform more safely than human drivers when compared in aggregate. In crafting our approach to this principle, we look to industry best practices in systems engineering methods,³⁰ coding and tool qualification standards,³¹ configuration management approaches,^{32,33} and safety culture models.³³ We also employ a variety of methods including simulation, track testing, and on-road testing to gauge system performance.

Proficient covers the following NHTSA safety elements: *System Safety*, *Operational Design Domain*, and *Object and Event Detection and Response*. □

The nominal operation of the self-driving system shall result in safer performance than human drivers.

- *Nominal operation* is performance in the absence of hardware and software faults, i.e. when everything is working as intended. We address what happens in the case of a detected fault in [section 06.02](#).
- We anticipate that demonstrating *safer performance than human drivers* will require that we quantify safe driving with tractable, credible metrics. Crash rates are one accepted measure, but are often subject to inconsistent reporting,^{25,26} miss important contextual factors,²⁷ and/or require an impractical magnitude of driving exposure.²⁸ Thus, we evaluate how our system performs over an aggregation of both common and rare scenarios, using measures that include traffic rule infractions and vehicle dynamics attributes.²⁹

²⁵ Farmer, 2003, 'Reliability of Police-Reported Information for Determining Crash and Injury Severity.'

²⁶ World Health Organization (WHO), 2009, 'Global Status Report on Road Safety: Time for Action.'

²⁷ Wang & Zhang, 2017, 'Analysis of Roadway and Environmental Factors Affecting Traffic Crash Severities.'

²⁸ RAND Corporation, 2016, 'Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability.'

²⁹ RAND Corporation, 2018, 'Measuring Automated Vehicle Safety: Forging a Framework.'

³⁰ Voirin, 2017, 'Model-based System and Architecture Engineering with the Arcadia Method.'

³¹ RTCA, 2011, 'DO-330 Software Tool Qualification Considerations.'

³² SAE International, 2011, 'EIA 649B Configuration Management Standard.'

³³ Institute of Electrical and Electronics Engineers (IEEE), 2012, 'IEEE 828 Standard for Configuration Management in Systems and Software Engineering.'

³⁴ National Aeronautics and Space Administration (NASA), 2015, 'NASA-HDBK-8709.24 NASA Safety Culture Handbook.'

System Safety

A robust systems engineering approach to ODD selection and characterization along with Object and Event Detection and Response (OEDR) serve as a crucial foundation to meeting this Safety Principle.

By limiting our self-driving vehicles to a specific ODD, we can mitigate risk from preventable harmful events. Additionally, by appropriately detecting and responding to actors and objects in the built and natural environment, we can ensure the self-driving system is operating and responding as intended to a variety of inputs.

It is through the combination of these two capability areas, evaluated by our rigorous verification and validation methods, that we can progress towards safer-than-human performance in a given ODD. □

Operational Design Domain

Before beginning any self-driving testing we establish the ODD. The ODD describes the specific conditions under which the self-driving system is intended to function, including where and when the system is designed to operate. This parameterization is not only designed to address the performance of the base vehicle platform but also system level capabilities, environmental scenarios, and appropriate self-driving system responses. We employ a three-step process to define the ODD: identify, characterize, and constrain.

Identifying the ODD

—


We begin by identifying specific geographies where we would like to ultimately deploy self-driving vehicles on the Uber network by taking into consideration a number of factors, including the regulatory environment, areas where we can extend our network's reach to better serve riders, and financial viability. Using the information layers of our high-definition maps, as well as data from Uber's core business, we convert the road geometries and static features of these geographies into a list of autonomy capability requirements for our self-driving vehicles. This list of requirements constitutes an intended production ODD.

This intended production ODD is converted into a technology roadmap, which describes the incremental expansion of our ODD to reflect new capabilities and the maturation of existing capabilities.

Uber's Self-Driving Safety Principles

06.01 → Proficient

Operational Design Domain (continued)



Characterizing the ODD

The ODD characterization process includes:

- **Driving the area manually** to collect detailed data and logs on the scenarios and actors that exist within the ODD.
- **Adding data tags** to camera and LIDAR footage collected from manually-driven logs, highlighting potentially relevant attributes of actors in and around the road as well as attributes of road design (e.g. road geometry or curvature, traffic control measures).
- **Synthesizing the tagged data** to identify and break down information on all scenarios and subsequent system behavior requirements for each scenario.
- **Creating representative simulation and track tests** to evaluate current and future software releases.

This process enables us to:

- Confirm requirements for self-driving system capabilities.
- Identify sufficient test coverage both through simulation and track testing to assess performance of the self-driving system before testing on public roads.
- Provide clear operational guidelines and performance requirements to support on-road operations, e.g. policies governing system takeovers and handling scenarios not captured in the pre-approved and established ODD.

Once we have characterized an ODD, the self-driving system must pass the identified set of offline tests and track tests before operating on public roads.

Operational Design Domain (continued)

Constraining the ODD

—

To prevent our self-driving system from operating outside of the intended ODD, we constrain the vehicle routing capability to only the approved ODD. We enforce these limitations using a Policy Constraints System. Policy Constraints restrict the routing of the self-driving vehicle based on a set of configurable ODD elements, e.g. road speed, road type, and traffic control devices.

Our Mission Specialists also monitor road conditions while operating in the field. Mission Specialists are trained on the governing ODD, and are prepared to take manual control of the vehicle when presented with a scenario that is not included in the current ODD. When one of our vehicles encounters a situation which the Mission Specialists know it is not yet capable of navigating in self-driving mode, e.g. a road blockage or closure, the left-seat Mission Specialist, or Pilot, is trained to take manual control of the vehicle and the right-seat Mission Specialist, the Co-Pilot, is trained to report the blockage. This report then initiates a process by which a constraint is generated to address the blockage and deployed throughout the fleet.


We mitigate risks posed by natural environmental factors during road operations by constraining driving to particular weather and road conditions. During development, local weather and events are assessed prior to deploying vehicles for on-road testing. If prevailing conditions are not in the vehicle's ODD, Mission Specialists are notified to disengage self-driving mode and/or cease further operations until it is safe to proceed, as covered in [section 06.02](#). □

Object and Event Detection and Response

Once the ODD is defined, we define and assess the appropriate system behaviors when detecting and responding to actors and scenarios in a given ODD. OEDR refers to the detection of any object or event that is relevant to the driving task, as well as the implementation of the appropriate response to such circumstances.³⁵ In order to ensure safe operation, the self-driving system must be capable of detecting and responding to a variety of static and dynamic objects in the road environment. The following sections describe how the self-driving software introduced in [section 05](#) delivers on this response.

³⁵ NHTSA, 2017, 'Automated Driving Systems 2.0: A Vision for Safety.'

Object and Event Detection and Response (continued)



Mapping

Understanding the Existing World

High-definition maps allow Uber's self-driving vehicle to understand the world in detail before it arrives at a particular location. By knowing and storing precise road information on a virtual map, the vehicle can anticipate proper behavior without requiring as much real-time scene understanding. Maps can improve safety by enabling the vehicle to anticipate the need to slow down or otherwise optimize its motion plan, e.g. before the Perception system is able to observe an upcoming tight turn.



Perception

Detecting the Environment, Actors, and Objects


As described in [section 05](#), our self-driving vehicles are equipped with a number of overlapping sensors gathering data covering 360° around the vehicle. Each sensing modality has its own strengths; combining these modalities provides a more complete, more accurate view of the environment.

Our Perception software detects and tracks individual actors and objects in order to generate estimates of their position and velocity and register other attributes that may inform their future motion. For example, turn signals and hazard lights may convey information about the intent of other vehicles. However, a car with its left turn signal on may not actually turn left, so, while the system perceives the turn signal, it continuously estimates position, orientation, speed, and other variables in order to ensure it can respond appropriately to the vehicle's ultimate course of action. The system also forms a view of stationary objects

Object and Event Detection and Response (continued)

that convey useful information that should govern its motion, e.g. reading the state of traffic lights.

The main detection and classification stages that operate on sensor data are machine learned modules that are trained and evaluated using extensive labeled datasets covering the ODD. These datasets are made more comprehensive and detailed over time through tooling, offline algorithms, and human efforts. In some cases, when an object or actor may not be properly classified, the system is designed to handle this class uncertainty. In addition to reasoning about uncertainty, the system has a second stage to account for actors or objects in the world that have sensor data but have not been explicitly detected as a known actor or class of object. For these cases, the system estimates the actor's extent and velocity and maintains a large amount of uncertainty in terms of future motion so the vehicle can react conservatively.



Prediction

Reasoning About What Actors And Objects Might Do

Our Perception software creates a representation of the driving environment and our Prediction software uses this representation to predict what the actors or objects in the environment are likely to do next. Some objects are fixed structures, such as buildings, ground, and vegetation, and we do not expect these objects to move. Actors, such as vehicles, pedestrians, bicyclists, and animals, are expected to move. Our software considers how and where all actors and objects may move over the next ten seconds.

Our prediction software applies different models of behavior for different actor and object classes: if an actor is perceived as a moving vehicle, it requires different possible predictions of, e.g. speed, direction, than if it were perceived as parked. If the Perception software is not able to positively confirm an actor's or object's classification or state of motion, it shares multiple potential options.

Object and Event Detection and Response (continued)

The Prediction software considers and presents multiple feasible object intents to the Motion Planning software, including intents which would put the actors or objects in the self-driving vehicle's path, even when the vehicle has the right-of-way. The Prediction system evaluates the probability that each behavioral model accurately describes what the actor or object is doing. The Motion Planning system uses these probabilities to effect an appropriate amount of caution in response to less predictable actors or objects. The system performs these predictions many times a second so as actors change direction or intent the system is designed to continually reassess their likely next move.



Routing, Navigation, and Motion Planning

Planning What To Do

Our Routing and Navigation software plans a route for the self-driving vehicle that takes it from its current position to its desired destination according to the rules of the road encoded in the map and any active constraints.

Our Motion Planning software combines information from the generated route, as well as perceived actors and objects and their anticipated movement from Perception and Prediction as inputs, and creates a motion plan for the vehicle.

Motion Planning provides for defined spatial buffers to be maintained at all times between the vehicle and other actors in the environment; the size of these buffers varies with speed. To preserve an appropriate buffer between the vehicle and any actors in the

Object and Event Detection and Response (continued)

environment, the system may opt to change lanes, brake to restore a safe following distance, or come to a controlled stop and wait until the situation clears.

Occlusions, or obstructed views, present significant challenges for both self-driving and human-driven vehicles. Our self-driving system reasons about occlusions and seeks to maintain the ability to avoid actors coming out of an occlusion at any reasonable speed. We have developed our system to be more conservative than typical human drivers with respect to occlusions.

Vehicle Control

Executing The Vehicle Plan

Vehicle Control executes the trajectory supplied by Motion Planning by controlling the actuation of the vehicle, including, e.g. steering, braking, turn signals, throttle, and gear, through communication interfaces. Further, vehicle control is responsible for understanding the dynamic limits and present condition of the vehicle, including any faults or error conditions that may affect Vehicle Control and communicates this information back to the self-driving sub-system.

We develop our Vehicle Control software in partnership with the manufacturer of the base vehicle, thereby ensuring the self-driving system understands the capabilities of the vehicle and is able to avoid conflict with base vehicle systems.

→ Looking Forward

Vehicle Control needs to provide highly-reliable operation, particularly in instances where the vehicle must be safely and immediately brought to a stop. Thus, we have chosen to develop Vehicle Control as a secondary computing system on embedded hardware that is distinct and independent from the self-driving computer.

This design provides fault tolerance through features such as redundancy, high integrity processors, and additional IMUs. This system is being developed taking into account best practice and industry standards for functional safety, including ISO 26262,³⁶ ISO 16750,³⁷ MISRA C 2012,³⁸ and AUTOSAR 4.2.³⁹

Uber is still developing a self-driving system that can safely operate without a human operator behind the wheel. As such, our system may not be capable of delivering on any specific driving behavior at present. In fact, we have frequently demonstrated proficiency on a specific scenario set only to identify a new variation beyond our current capability. Uber is purposely not including a list of behavior competencies in this report: we believe even the behaviors which we have routinely found our system capable of handling with no operator intervention require more testing, more variations, and potentially more development. □

³⁶ ISO, 2011, 'ISO 26262 Functional Safety for Road Vehicles.'

³⁷ ISO, 2012, 'ISO 16750-2:2012.'

³⁸ Motor Industry Software Reliability Association (MISRA), 2013, 'MISRA C:2012.'

³⁹ Automotive Open System Architecture (AUTOSAR), 2013, 'AUTOSAR Classic Platform Release 4.2.'

Mission Specialists

Mission Specialists play an essential role in the safe development of self-driving vehicles by enabling greater collaboration between software, hardware, and test teams. While we are still in the developmental phase, Mission Specialists are the ‘humans in the loop,’ should the system be unable to maintain control. The ability of our Mission Specialists to maintain control of the vehicle during testing is addressed through self-driving system design, training, and operational policies.

Today, we operate our self-driving vehicles with two Mission Specialists in the vehicle. The Pilot, or operator behind the steering wheel, is solely focused on ensuring safe operation of the vehicle, while the Co-Pilot, the second operator in the right front seat, is tasked with monitoring and annotating the behavior of the self-driving system via a laptop. We previously operated a portion of our fleet with a single operator behind the wheel and no Co-Pilot. We believe that operating with two Mission Specialists reduces workload and potential for fatigue, distraction, or misuse.

Our Mission Specialists are key to understanding and evaluating the performance of our self-driving system. They bring the evolution of feature development full circle by providing significant insights from offline, track, and road testing. Proper training, continuous education, and open lines of communication back to our engineering teams ensure they are able to do their jobs safely, effectively, and efficiently. We believe best practices for safe self-driving operation should continue to develop through an open discussion among self-driving developers.

Hiring and Screening

Because of the elevated operational responsibilities of our Mission Specialists relative to standard driver-partners, candidates undergo a multi-step interview process, which assesses technical, communication, and reasoning skills, in addition to physical vehicle control.

1. Application Review and Phone Screen

The Mission Specialist hiring process begins with an application review and phone screen conducted by our recruiting team. Recruiters screen for technical competency and testing experience, safety awareness and training, and driving history through a series of standardized screening questions. Once the recruiting team completes the screening report, hiring managers determine if the candidate meets minimum requirements to move to the next stage.

2. Homework

Candidates are asked to complete a homework assignment which involves identifying the navigation path for a self-driving vehicle in a common traffic scenario. This homework assignment is used to identify the candidate's ability to detail an otherwise nominal traffic situation in a way which would help a developer model software around the specific scenario. Strong candidates offer multiple solutions and account for constraints (e.g. environmental and safety), present information concisely and completely, and demonstrate the ability to research and synthesize information.

3. Onsite Interview

Candidates are invited to in-person interviews with hiring managers who assess their understanding of the position and qualifications. Managers ask standardized questions to assess the candidate's


Mission Specialists (continued)

competency for safety procedures and their ability to work through difficult situational scenarios. The in-person interview consists of an in-vehicle driving and technical skills evaluation. Candidates are assessed on their ability to safely and responsibly operate a vehicle in manual, multitask, and effectively provide information about the driving environment and technology.

4. Debrief and Hiring Decision

Hiring managers and interviewers work with recruiting teams to determine if the candidate exceeds requirements set for the position. All candidates are also subject to certain screenings including a motor vehicle record check.

Prior to operating a self-driving vehicle, Mission Specialists undergo extensive training on our self-driving vehicles including the software, hardware, and operating skills. We believe this to be a critical fail-safe for our developmental self-driving system. Mission Specialist training is covered in [section 06.02](#). □



Uber's Self-Driving Safety Principles

06.02 → Fail-Safe

Principle 2 Fail-Safe



In addition to demonstrating that our system is safe when it is working correctly, we also have to demonstrate that it is safe when it encounters a fault.

To fulfill this principle, we partition safety responsibilities to different parts of the system; we also institute fallback maneuvers during system-level failures. Any part of the vehicle – base vehicle components, add-on electronics, or our software – has the potential to experience a failure during operation. We contain these risks by minimizing common-cause failures through system architectural analysis.

Fail-Safety covers the following NHTSA safety elements: *System Safety and Fallback (Minimal Risk Condition)*. □

Any safety-relevant failure shall result in transition of the vehicle to a minimal risk condition or shall be extremely improbable.

- A *safety-relevant failure* is a malfunction that results in reasonable probability of harm to a person. Other types of failures may result in non-safety related outcomes, e.g. a poor experience for a rider.
- The *minimal risk condition* is a system state which “reduce[s] the risk of a crash when a given trip cannot or should not be completed... It may entail automatically bringing the vehicle to a stop within its current travel path, or it may entail a more extensive maneuver designed to remove the vehicle from an active lane of traffic and/or to automatically return the vehicle to a dispatching facility.”⁴⁰ The appropriate maneuver depends on the particulars of the failure and the circumstances of the scenario.
- While we seek to eliminate all unguarded failures, we do allow for the possibility that some could persist, if and only if it can be ensured that their probability of occurrence is exceedingly remote and/or the potential severity is limited. Permitting an *extremely improbable* safety-relevant failure to persist borrows from aviation risk frameworks.⁴¹

⁴⁰ SAE International, 2018, ‘J3016_201806: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles.’

⁴¹ Federal Aviation Administration, 2011, ‘System Safety Analysis and Assessment for Part 23 Airplanes.’

System Safety

No system is immune from conditions that interfere with its ability to correctly execute its intended function. These conditions are faults; the related loss of functionality is a failure.⁴² Interruptions to self-driving system functionality, without appropriate mitigations, can pose a risk to the safety of those in and near the vehicle. A self-driving vehicle should therefore be designed, to the extent practicable, to function predictably, controllably, and safely in the presence of any faults and failures.

In mitigating a safety-relevant failure, we must make it extremely improbable, prevent it, and/or implement a solution that transitions the vehicle to a minimal risk condition in the case of a failure. We do this by using robust and thoroughly tested components, designing key redundancies into the system, and implementing software that monitors the system for faults and takes action when they occur. Redundancies and fault detection software are tenets of fault tolerant system design. □

Fallback

Minimal Risk Condition

Preventing Faults

— Our approach to fault prevention is informed by similar approaches in other industries, such as automotive and aerospace. For example, we leverage processes from ISO 26262,⁴³ an automotive industry standard, to identify, assess, and mitigate faults and hazards for electrical and electronic components.

Designing a Fault-Tolerant System

— Self-driving vehicles must be able to tolerate faults. Fault tolerance requires that the self-driving system is able to retain certain functionality even when faults occur. When faced with a safety-relevant fault, the system can either return control to the Mission Specialist, immediately bring the vehicle to a safe stop, or pull over when safe to do so. Today, we rely on Mission Specialists to resume control of the vehicle in the presence of a safety-relevant fault by alerting them to a transition out of self-driving mode via audio and visual cues. Transitions into and out of self-driving mode are covered in [section 06.04](#).

System-level fault protection involves implementing mitigations that transition the vehicle to a minimal risk condition in the case of a safety-relevant failure. The self-driving vehicle is being designed to detect that a fault has occurred and initiate a fail-safe, or fallback response. The fault management system must discern potential impact from individual and aggregate faults, prioritize the most potentially harmful, cascade related dependencies, and transition the system to a minimal risk condition.

⁴² Consistent with failure (1.39) and fault (1.42) definitions in ISO, 2011, 'ISO 26262 Functional Safety for Road Vehicles.'

⁴³ ISO, 2011, 'ISO 26262 Functional Safety for Road Vehicles.'

Uber’s Self-Driving Safety Principles

06.02 → Fail-Safe

Fallback Minimal Risk Condition (continued)

The system determines its response to a safety-relevant fault based on two factors:

1. The functionality, if any, that the system retains in the presence of the fault.
2. The time it takes from the occurrence of the fault until a harmful event could occur as a result.

→ **Looking Forward**

We intend our next generation vehicles will feature additional redundancies, such as redundant steering, braking, and immobilization systems. □

Examples of Potential Fault Mitigations a Self-Driving Vehicle May Utilize

Fault	Fault Type	Mitigation Plan
Primary Compute Power Failure	Electrical Power Systems	Backup power turns on, the system detects the fault, and the vehicle is safely brought to a stop.
Loss of Primary Compute or Motion Planner Timeout	Self-Driving Software	If the VIM stops receiving trajectories from the self-driving system, the Mission Specialist will be notified via LED status lights and an audio cue that the vehicle has returned to manual mode. In our next generation vehicle, the VIM will bring the self-driving vehicle to a safe stop along the most recently received valid route, while using IMUs and wheel speed sensors to maintain control of the vehicle.
Wheel Speed Sensor Data Delay	Vehicle Platform	Our systems monitor the data coming from the vehicle's wheel speed sensors and will detect if the data becomes delayed or stops being sent. The system will then initiate a safe stop.
Door Opens While Driving at Speed	Misuse	Our systems will detect that the door has opened and will safely stop the vehicle.

Training for a Fault-Tolerant System

In some scenarios, the most appropriate fallback response is to return control of the vehicle to the Mission Specialist. Mission Specialists are required to complete a comprehensive training program which prepares them to safely operate a self-driving vehicle and protect its equipment from damage.

Manual Driving

Every Mission Specialist must be capable of safely operating our vehicles, whether in manual or self-driving mode. For this reason, we open our training program with safe manual driving habits, first on a closed course, followed by public road training.

Driving Training

- **Driving dynamics and awareness** via in-classroom instruction and driving drills on the test track.
- **Emergency maneuver exercises**, including collision avoidance, anti-lock braking, and slalom driving at speeds, as relevant to our ODD.
- **Parking and reversing exercises** to assess spatial awareness, vehicle size limitations, vehicle placement relative to other actors, and the proper use of mirrors and reverse cameras.
- **Navigating occluded views** during manual driving.
- **Defensive driving**⁴⁴ online course to educate Mission Specialists on poor driving habits and new defensive driving techniques for operating self-driving vehicles.

ODD and Vehicle Platform Training

- **Overview of ODD** to educate Mission Specialists on its scope and required capabilities.
- **Overview of traffic laws** relevant to the ODD.
- **Incident response simulation** to practice confident handling of incidents.
- **Platform failures exposure** and assessment.
- **Volvo Advanced Driver Assistance System** explanation.

Technical Education

To safely operate a self-driving vehicle, a Mission Specialist must understand the essentials of the self-driving computer. Mission Specialists undergo extensive software and hardware training on:

- **The Self-Driving Software**
The training program uses the self-driving system architecture to explain how the vehicle makes decisions. This includes a thorough review of maps, sensors, Localization, Perception, Prediction, Routing and Navigation, Motion Planning, and Vehicle Control.
- **The Hardware on the Vehicle**
The program reviews sensor functions and limitations, as well as vehicle control hardware. Mission Specialist trainers also demonstrate radar range, vehicle positioning, LIDAR blind spots, and camera angles and views.

⁴⁴ NSC, 2017, 'Defensive Driving Online Course - 4 Hours.'

Training for a Fault-Tolerant System (continued)

Our training program includes modules on vehicle capabilities, limitations of the hardware and software, and how the self-driving vehicle reasons about and interacts with its environment. Training modules include:

- **Software limitations** describes the vehicle's capabilities in the ODD.
- **Occluded views module** explains how self-driving vehicles identify and handle occlusion. The content in this section covers the vehicle's capabilities in managing occlusion and proper procedures for piloting a self-driving vehicle through an occluded intersection. Exercises are both in-classroom and in-vehicle.
- **Pedestrians and cyclists interactions** demonstrates how the self-driving vehicle responds to pedestrians and cyclists. Exercises are both in-classroom and in-vehicle.

Piloting

—

Before operating a vehicle in self-driving mode, a Mission Specialist must complete piloting training. As with the manual driver training, we first introduce these fundamentals in the classroom and on a closed-course track prior to public roads.

Piloting Fundamentals

- **Engaging and disengaging techniques** cover procedures on how to safely engage and disengage self-driving mode, first in a stationary vehicle then in a moving vehicle. This course also covers the vehicle controls and nominal self-driving operations as well as the visual and audio cues that are

presented upon system state transition. For more on transitioning between manual and self-driving mode, see [section 06.04](#).

- **Safety and personalization** covers adjusting mirrors and seat position to properly and comfortably pilot the vehicle.
- **Touch grip** training covers proper hand position on the steering wheel. This hand position allows the Mission Specialist to disengage from self-driving mode using the steering wheel when appropriate.
- **Pedal Shadowing** covers disengagement from self-driving mode by depressing the accelerator or brake pedals. Mission Specialist are trained to hover a foot over the proper pedal to ensure a smooth and safe takeover if the vehicle is in motion.
- **Front Seat Control App (FSCA) interactions** cover the policies for interacting with the touchscreen.

Fault Injection Training

During Fault Injection Training (FIT), trainers inject faults into the system so trainees can safely gain exposure to the vehicle's capabilities in a variety of fault situations. This module takes place on a test track and has three parts:

- **Basic FIT** exposes trainees to in-vehicle faults without the added complexity of environmental factors or outside actors and establishes a baseline reaction time for the trainee up to the maximum system capability, independent of context or environment. This module covers correct mechanics such as touch grip and pedal shadowing, vehicle controllers, and scenarios that can lead to faults or situations where these faults may become problematic.

Training for a Fault-Tolerant System (continued)

- **Self-Driving FIT** exposes trainees to the faults covered in the basic FIT module with the addition of environmental factors. For example, trainees will experience faults in intersections, before intersections, just after intersections, in turns, and on straightaways.
- **ODD Scenario FIT** focuses on environmental faults or actors and scenarios that are currently out of scope for the self-driving system and will therefore require that the Mission Specialist take over, e.g. lane blockages.

Co-Piloting

– We operate with a Mission Specialist in both the front left and right seats of the vehicle today. Mission Specialists are trained how to Co-Pilot in the right seat, including training on the Operator Control Station (OCS) laptop and effective communication of complex technical issues.

Continuous Education

– During development, vehicle behaviors, capabilities, and system-level features are constantly evolving. In order to provide our Mission Specialists with the most up-to-date information on our system, we hold daily mission briefings and require completion of online learning modules and in-vehicle or in-classroom training. We intend to have Mission Specialists train in new capabilities and functionalities before operating in those capabilities or functionalities; examples of changes resulting in new training include enabling self-driving lane changes, increasing vehicle operating speeds, and expanding ODDs. □

Operational Safety

To ensure a high level of proficiency in day-to-day operations, Mission Specialists must be aware of and responsive to their operating environment, both inside and outside of the vehicle.

Understanding the ODD

– As described, we train Mission Specialists in the classroom on the ODD, including the limits of the self-driving system, how and when to resume manual control of the vehicle, proactively or in the event of a system fault or failure. This information is reinforced during in-vehicle training and FIT training modules. This training prepares the Co-Pilot to inform the Pilot of any upcoming events that may require a transition to manual mode. For more on transitioning between driving modes, see [section 06.04](#).

Mission Specialists also receive a daily, pre-mission briefing on the current operational test plan, ODD, and software release status. As the ODD evolves, we brief or train Mission Specialists, depending on the scope of the change.

Communication

– Effective communication between the Pilot and Co-Pilot plays an important role in safe self-driving vehicle operations. Our training program covers guidelines for managing in-vehicle communication. Further, Mission Specialists are trained to communicate relevant information from the OCS that can assist the Pilot.

Operational Safety (continued)

Fatigued / Distracted Driving Prevention

—

In light of the unique opportunities for distraction and fatigue while operating self-driving vehicles, our training programs focus on assisting Mission Specialists in recognizing and managing these situations.

- Our *Distracted Driving* module raises awareness of distracted driving, possible consequences, and steps to avoid this behavior. Mission Specialists read and discuss the National Safety Council's (NSC's) "Understanding the Distracted Brain"⁴⁵ and complete exercises to ground their learning.
- Our *Fatigued Driving Prevention* module references guidance from the U.S. National Transportation Safety Board (NTSB)⁴⁶ and U.S. Federal Motor Carrier Safety Administration (FMCSA).⁴⁷

For more on our preventative approach to distracted driving and fatigue, see [section 06.04](#).

As the performance of the self-driving system increases, Mission Specialists may become less effective as the frequency of intervention decreases. We appreciate the importance of mitigating this risk and intend to continue to undertake studies on human factors, effective assistive measures, and overall support structures for safe, self-driving vehicle operations. □

⁴⁵ NSC, 2012, 'Understanding the distracted brain.'

⁴⁶ NTSB, 2017, 'NTSB 2017-2018 Most Wanted List of Transportation Safety Improvements: Reduce Fatigue-Related Accidents.'

⁴⁷ FMCSA, 2014, 'CMV Driving Tips - Driver Fatigue.'

Uber's Self-Driving Safety Principles

06.03 → Continuously Improving

Principle 3 Continuously Improving



To fulfill this principle, we draw on quality processes for software development and hardware component production. We have implemented and refined workflows for collecting and analyzing test results from both offline and track testing, as well as on-road driving. From concept design to public road testing, our vehicles and self-driving system components pass through manufacturing and commissioning tests. We run a series of standardized tests on each software release, and leverage standardized documentation and issue tracking tools to effectively capture learnings and see them through to a resolution. We currently rely on Mission Specialists' feedback in addition to automated results.

Continuously Improving covers the following NHTSA safety elements: *Validation Methods*. □

Any observed anomalies shall be systematically reported, evaluated, and resolved with appropriate corrective and preventative actions.

- An *anomaly* is an undesirable and unexpected behavior or result.⁴⁸ We are attentive to these kinds of events as warning signs of potential safety issues before they result in harm.
- We implement processes and mechanisms to *consistently capture and assess* the severity of observed issues so that we can assess the potential impact of these issues on continued safe operation.
- In response to an identified anomaly, we *determine and execute an appropriate action*. This may include, e.g. implementing a hardware or software fix, changing operational procedures temporarily or permanently, or determining that, while unexpected, the observance does not indicate an underlying safety risk.

⁴⁸ Consistent with anomaly (1.2) definition in ISO, 2011, '[ISO 26262 Functional Safety for Road Vehicles](#).'

06.03 → Continuously Improving

Uber's Self-Driving Safety Principles

Validation Methods

We have built robust tools and processes at every step in our development cycle to track and respond to system issues. Both our software and hardware development processes thoroughly evaluate the self-driving system prior to any testing on public roads. □

Self-Driving Software Quality Processes

Uber employs a rigorous testing and validation process from an initial software change through real-world testing.

Deviations from expected operation during offline, test track, and on-road testing and data collection are recorded and shared with self-driving system development teams. In particular, comments from our Mission Specialists provide a firsthand account of the in-vehicle rider experience. Data created by noteworthy events, such as large deviations from one planned trajectory to the next, system diagnostics, or Mission Specialist interventions, are identified for our data review process.


We then track these events from discovery to resolution. For example, we may re-simulate software changes to evaluate their impact on these key events, to confirm failures are resolved as intended, and are not reintroduced. Many events are incorporated into datasets for machine-learning algorithms, while others are utilized as challenging test cases. Software failure events are also replicated in simulation scenarios, which can be varied. Once issues are resolved, the resolution factors into every new software change with any eye to preventing the accidental re-introduction of previous undesirable behaviors.

Uber's Self-Driving Safety Principles

06.03 → Continuously Improving

Self-Driving Software Quality Processes (continued)

Software Development and Validation Process



Offline Testing

We have developed a suite of offline testing tools that enable us to test code as soon as it is written, providing valuable insight into potential issues as early as possible in the development lifecycle.

Each software release is subjected to a battery of automated offline tests that provide a baseline level of confidence that the release should be viable for advancing to additional release testing stages. If a release does not pass the offline release evaluation process, it does not move forward. A sample of offline release tests include:

- **Map Compatibility Test**
As our self-driving software requires a map in order to support autonomous operations, this test ensures both the latest map and self-driving software release being tested have no integration issues.

- **Onboard Integration Tests**
This set of tests confirms that the latest self-driving software release connects to the vehicle platform as desired, notifications are passed correctly between the software, self-driving system hardware, and base vehicle, and the software release has been correctly deployed to the vehicle.
- **Unit Tests**
Tests that are designed to test atomic (non-divisible) portions of code, and are run independently on software changes prior to landing on the code base.
- **Virtual Simulation Regression Set Test**
Set of simulations representative of nominal on-road scenarios against which all software releases are tested for regression, i.e. when the simulated self-driving system behavior fails a scenario that it previously had passed.
- **Reaction Time Metrics Test**
Evaluates whether the reaction time of the self-driving system software meets our expected requirements.

Uber's Self-Driving Safety Principles

06.03 → Continuously Improving

Self-Driving Software Quality Processes (continued)

We handle failures in relationship to where they arise in the software release process:

- **Failures in Testing Prior to Integration**
Failures found prior to integration into the master code repository are generally provided directly to the author of the code change for resolution as part of the peer review and development testing process.
- **Failures in Release Testing**
We treat failures found during release testing phases as very high priority, because a failure in release testing indicates the presence of a flaw in the codebase. We conduct an initial triage of the failure to understand root cause to support a resolution.
- **Preserving Accuracy of the Testing Regime**
In addition to addressing anomalies in the system, we work to prevent and address anomalies in the testing regime itself. When a test has been flagged as unsound, meaning failures could possibly be a false-positive, the test owner reviews the test for validity and/or revision.

Where possible, in addition to tracking the initial failure, a parallel effort begins to create an offline test that seeks to detect this failure prior to entering the codebase in the future. This ensures that our battery of tests becomes more comprehensive over time. Additionally, if an issue is identified as safety critical, we have the option to immediately and remotely stop operations entirely or within a specific ODD through real-time communications with the vehicle fleet.

Hardware in the Loop (HIL) Testing

HIL testing is concerned with ensuring performance of our software when running on representative hardware. By coupling our self-driving system

software with our self-driving system hardware prior to it actually being placed on a vehicle, we are able to isolate and diagnose faults that could not be revealed through software testing alone. HIL testing is required for many embedded software changes prior to releasing on a vehicle.

Simulation

Simulation plays a key role in self-driving software development: it enables testing of relatively rare, challenging scenarios without the physical risk associated with test track or on-road testing, and it also allows testing more routine scenarios with minor, controlled variations. Simulation tests have different permutations and combinations of traffic patterns, speeds, and trajectories for all the actors and objects in a scenario, including our self-driving vehicle.

Benefits of simulated driving test approaches include:

- **Safety**
Simulations allow us to test high risk scenarios safely that would be dangerous to test in the real world.
- **Repeatability**
Simulations can be rerun in the same exact way over time. This predictable deterministic setup allows us to evaluate progress of subsequent builds of self-driving software against the same scenarios with a degree of repeatability that is not possible by track testing.
- **Frequency of occurrence**
Many of the challenging scenarios we need to test occur infrequently in the real world. In simulation, we can increase the frequency of these scenarios in order to test our systems' ability to handle lower-probability events.

Self-Driving Software Quality Processes (continued)

- **Variance**
We can run numerous variations of the same test scenario.
- **Efficiency**
It is safer and more economical to stress test our self-driving system in simulation than on our test track or on public roads.

We are focused on ensuring reliability of simulation results and on measuring and improving consistency between our real-world and simulated vehicles, and between our test track and simulated scenarios.

Test Scenario Development

Fundamental to Uber's strategy for the development of safe self-driving technology is alignment between design, test, and use. We employ a scenario and ODD development framework that characterizes design requirements, real-world events (such as those collected through driving logs), synthetic test scenarios, and operational policies using a unified schema.

A scenario includes the physical environment as well as actors or objects and their static or dynamic paths. Each scenario is defined by a number of criteria for success, including, e.g. speeds, distances, and descriptions of safe behavior. Our scenario documentation provides the basis for virtual scenario builds that can be run in simulation and on our test track. Scenario success criteria are aligned with applicable traffic laws.

Our testing battery includes virtual models of scenarios that:

- **Require basic driving skills**
We identify and define basic driving capabilities

necessary to operate in a given ODD during the ODD characterization phase as outlined in [section 06.01](#).

- **Are likely to lead to crashes**
We are developing a set of scenarios that typically lead to crashes, based on an assessment of our own data and frameworks from NHTSA,⁴⁹ PROSPECT Project⁵⁰ and The European New Car Assessment Programme (Euro NCAP).⁵¹
- **Are particularly challenging for self-driving vehicles**
We add additional scenarios as they are identified through on-road operations or observed during offline testing.
- **Are ultimately intended to be representative of everything our vehicles could encounter in the real world**
The world can create an infinite number of unique cases. Human drivers can reason the correct action even in scenarios never encountered. Our self-driving vehicles should do the same. Our goal is to create a set of scenarios that represent our ODD. As we encounter new scenarios that are not covered, we intend to add or substitute scenarios to improve the set.

Track Validation Testing

— Software releases which have passed their offline testing advance to Track Verification Testing (TVT). We test and validate each software release on our closed course test track by subjecting the software to an appropriate set of fault-injected, performance-based, field-derived tests.

⁴⁹ NHTSA, 2007, 'Pre-Crash Scenario Typology for Crash Avoidance Research.'

⁵⁰ PROSPECT Project, ND, 'PROSPECT Project.'

⁵¹ Euro NCAP, 2018, 'Vulnerable Road User Protection.'

Uber's Self-Driving Safety Principles

06.03 → Continuously Improving

Self-Driving Software Quality Processes (continued)

- **Track Test Development**

We develop the test suite for TVT through an iterative process that begins with identifying our target ODD, understanding the capabilities required in that ODD, and developing tests to measure performance of the system. TVT is comprised of both on-vehicle tests, which exercise the self-driving vehicle's behavior, and offboard functionality tests, e.g. operability on the Uber network, all within the controlled ODD of the test track.

- **Analysis**

TVT is conducted multiple times over the course of a week to evaluate self-driving system performance with clearly defined pass/fail criteria. Criteria are established in capability-based product requirement documents, test plan and procedure documents, and design specifications.

If the system fails a test, the basis for that determination is documented and tracked using a standardized system.

- **Release Reporting**

For each software release tested through TVT, we generate a report to show performance of the self-driving system, including pass/fail percentage across all tests and breakdowns of problems encountered. After identifying and characterizing an ODD and having demonstrated proficiency against a set of representative set of offline and track tests, a self-driving software release is deemed ready for on-road operation.

On-Road Testing

– We believe that the potential of self-driving vehicles will only be realized if we are able to learn from real-world situations, while gaining and preserving public trust. On-road driving allows us to observe - in real time - the performance of our system when faced with the diverse set of inputs that cannot be fully anticipated or replicated in artificial environments; this controlled exposure under the supervision of our Mission Specialists enables us to both improve our technology in response to observed events as well as to prudently augment our virtual world and test track scenarios for greater test coverage on future releases. □

Self-Driving System Quality Processes

We also look to address potential hardware and software issues with the base vehicle platform and/or the self-driving system hardware via quality processes.

Design Quality

Component Level Design Verification

Hardware modules for our self-driving system undergo testing to confirm they are functioning properly, and to identify performance limits. Once we have confirmed nominal function, individual components are validated via environmental qualification testing. This testing provides comprehensive coverage of thermal, vibrational, electromagnetic and other environmental factors beyond what is expected during normal operation. In addition, components undergo extensive reliability testing to ensure proper functionality throughout the intended product lifecycle. This testing exposes components to wear and tear, namely to simulate lifetime exposure and ensure no degradation of function or performance.

Subsystem-Level Design Verification

We test certain subsystems in order to confirm effective interactions between components. This stage of testing involves HIL and simulation testing across hardware and software interfaces in a controlled environment. We also perform fault injection testing at this level. Automation of tests makes it possible

to conduct highly repeatable structured testing of hardware/software interfaces. This subsystem-level design verification is required before track and road testing.

System-Level Design Verification

Self-driving hardware and software components are integrated into the vehicle and tested to confirm performance of:

- **Mechanical interfaces** including thermal and structural integration into the base vehicle.
- **Electrical interfaces** including integration into the base vehicle power distribution system and onboard communication busses.
- **Control path interfaces** including Application Programming Interfaces (APIs) to provide base vehicle platform motion (e.g. steering, braking and acceleration) as well as other key actuations (e.g. turn signals and gear changes).

We confirm the Vehicle Control path through structured track testing focused on the system's ability to maintain control of the vehicle through a full range of maneuvers while testing other factors that are difficult to simulate. Results from sub-system HIL and simulation testing are confirmed with on-road and in-vehicle testing.

Uber's Self-Driving Safety Principles

06.03 → Continuously Improving

Self-Driving System Quality Processes (continued)

Manufacturing Quality

Uber has implemented a comprehensive set of quality control processes. We follow an internal process derived from the principles of ISO 9001⁵² for both assemblies built in-house and sub-systems received from suppliers. Our process is described in the list below.

- **Supplier Selection**

Supplier selection is conducted relative to the part or component being sourced, typically through an RFX process or approved vendor list.

As development of a build or module matures, we develop a quality control plan jointly with the supplier which specifies the type and frequency of quality data recording. The quality control plan reflects the complexity of the product, maturity/stability of the process, and statistical significance of the sample size.

- **First Article Inspection Process**

At the start of manufacturing, the supplier produces a small first batch of parts which is subjected to detailed inspection against specifications. The first articles are inspected to the design data package and approved/rejected by hardware design, manufacturing, and quality engineering teams. The first article inspection is required to authorize manufacturing of larger quantities and serves as a trial run for quality data recording.

- **In-Process Inspection Plan**

The inspection plan is created before the assemblies are built based on the design data package. The plan informs the production technician team of pass/fail criteria for component assembly.

We implement in-process quality checks at assembly stations; these must pass before moving the assembly to the next station.

- **Traceability**

All assemblies built in-house or from suppliers require traceability data recording of date/lot codes, component/module serial numbers, and revision tracking.

Uber assembly stations are set up with fastener tightening data recording. Fastener tightening data is collected through the use of smart tools. All torque tools and other applicable assembly tools are periodically calibrated and we retain records for the life of the tool.

- **End-of-Manufacturing-Line Testing and Outgoing Quality Control**

The upfitted vehicle undergoes software updates and an extensive series of tests to ensure hardware performance when operating as a system. We develop the end-of-line testing plan in close coordination with design engineering. We undertake ongoing inspection, including redundant checks of critical fasteners, fit and finish checks, review of traceability data, and documentation to create the vehicle assembly quality data package. We retain records for the service life of the as-built vehicle.

- **Calibration**

After passing outgoing quality control, the upfitted vehicle leaves the manufacturing facility and advances to calibration, road-released software loading, and closed course testing before on-road testing.

⁵² ISO, 2015, 'ISO 9001 Quality Management Systems.'

Uber's Self-Driving Safety Principles

06.03 → Continuously Improving

Self-Driving System Quality Processes (continued)

Operational Quality

Commissioning and Calibration

Commissioning and calibration is the final phase of the self-driving system quality process. The purpose of this phase is to ensure all sensors required for self-driving capabilities are fully functional, and to collect the data necessary to perform intrinsic and extrinsic calibration, or measurement of hardware parameters required to align and combine the data produced by multiple sensors. We operate the vehicle in order to expose it to specific targets and environments. Data logs are collected and we undertake quality assessment. Finally, we put the vehicle through a final driving test. Any issues identified throughout the process are tracked and resolved through a ticket system by trained technicians.

Maintenance and Repair

Uber-managed self-driving vehicle fleets undergo extensive maintenance and monitoring routines to ensure they continue to perform as expected. Prior to performing a day's mission, Mission Specialists subject the self-driving system to health checks and inspections to ensure it is ready for operations. We track pre- and post-operational inspections digitally, and the results automatically generate issues tickets for tracking. This ensures every vehicle is properly inspected before it leaves our testing operations centers.

We track all platform software and hardware issues that emerge during commissioning and operations through to resolution. Once an issue is identified, our trained vehicle technicians are responsible for verifying, analyzing, isolating, repairing, and confirming operational viability across all Uber's self-driving vehicles. □

Internal Safety Concern Reporting System

As part of continuously improving the way Uber develops self-driving vehicles, we have implemented an internal anonymous safety concern reporting system designed to collect valuable feedback from anyone on our team. We openly encourage our employees to raise awareness of any concern that, if addressed, has the potential to improve the safety of our self-driving operations.

Voluntary reporting systems have been successful in similar industries such as aviation⁵³ and health care.⁵⁴ Concerns can be reported from a named party or anonymously, to further remove possible disincentives to report. Concerns can also be reported directly to our team or, in the alternative, to company personnel outside the Uber ATG reporting structure. No punitive action will be taken against the reporter simply for the fact of lodging a safety concern. Each concern is taken seriously and assessed for potential safety risk, analyzed, reviewed, and resolved with appropriate corrective actions. We periodically highlight the reporting system internally, continuing to raise awareness around a proactive safety culture. □

⁵³ FAA, 2011, 'AC 00-46E - Aviation Safety Reporting Program.'

⁵⁴ Grant & Larson, 2007, 'Effect of an anonymous reporting system on near-miss and harmful medical error reporting in a pediatric intensive care unit.'

Uber's Self-Driving Safety Principles

06.04 → Resilient

Principle 4 Resilient



Our self-driving vehicles will not operate in a vacuum. They will encounter all types of road users, including other vehicles, pedestrians, bicyclists, scooters, and more; pick up and transport riders; and serve as a potential target for people with illicit motives. We must consider the ways one of our self-driving vehicles might be used or interacted with differently than intended and put in place reasonable protections. These behaviors may not be frequently expected and they may be intentional or unintentional.

Misuse scenarios undergo a risk analysis to determine the likelihood of occurrence and severity of the outcome(s). Reasonably foreseeable misuse focuses on human behavior, which cannot be fully

characterized or controlled. For this reason, we have tailored the risk schema from ISO 26262⁵⁷ and the U.S. Department of Defense (DOD) MIL-STD-882.⁵⁸ High-risk scenarios are those scenarios that are very likely to occur and result in a high-severity outcome.

To fulfill this principle, we undertake a systematic process to:

- **Identify potential sources of misuse, from riders to other road users, to would-be cyber intruders, and generate misuse scenarios.**

We undertake various research efforts to generate misuse scenarios. Drawing on this research, we define actors and vehicle 'moments,' or points in

Potential harm from reasonably foreseeable misuse and other unavoidable events shall be mitigated.

- We anticipate *reasonably foreseeable misuse* - scenarios in which our technology is used counter to its design or purpose - because self-driving vehicles, like any other technology, are subject to an innumerable set of theoretical misuse scenarios.⁵⁵
- *Mitigation* in the context of this misuse involves preventing, protecting, and/or warning against potential harm; steps should be undertaken as possible in that order.⁵⁶
- Types of *misuse* considered under this principle are remote threats or malicious access to our self-driving computer.
- There may be situations where a crash is *unavoidable*, or beyond our control, due to the actions of other road users. In this case, we will work to minimize the likelihood and severity of harm.

⁵⁵ This is consistent with the retiring of extremely unusual scenarios that is permitted in automotive hazard analysis per Clause 7 of ISO 26262-3:2011, which gives as an example the scenario of a vehicle involved in an incident which includes an aeroplane landing on a highway (see Annex B.3). See ISO, 2011, '[ISO 26262 Functional Safety for Road Vehicles.](#)'

⁵⁶ This is consistent with guidance of §174 of the European Commission's (EC's) Guide to Application of the Machinery Directive 2006/42/EC. See EC European Agency for Safety and Health at Work, 2010, '[Guide to application of the Machinery Directive 2006/42/EC.](#)'

⁵⁷ ISO, 2011, '[ISO 26262 Functional Safety for Road Vehicles.](#)'

⁵⁸ U.S. DOD, 2012, '[MIL-STD-882E System Safety.](#)'

Principle 4 Resilient (continued)

time, during the vehicle's lifecycle, e.g. picking up or dropping off riders. We envision that the number of misuse scenarios will grow over time, so this analysis is continuous. However, by defining the actors and moments, we can take a systematic approach to defining the different permutations of interactions.

- **Assess the inherent risk, identify potential mitigations, and validate the effectiveness of our mitigations.**

Once a misuse scenario has gone through risk analysis, we design and implement an appropriate mitigation to minimize the likelihood of the misuse and severity of impact. In cases where risk cannot be eliminated completely, we aim to prevent a severe outcome, deter harmful human behavior, and/or put in place clear response policies to reduce impact. All mitigations go through a verification and validation process, as described in [section 06.03](#). The data-gathering process, including on-road testing, is aimed at continuously improving our ability to identify and respond to these types of scenarios.

Resilient covers the following NHTSA safety elements: *Human-Machine Interface, Crashworthiness, Vehicle Cybersecurity, and Data Recording*. □

Human-Machine Interface

Mission Specialists' Behaviors

– Human drivers are constantly receiving new information from the driving environment, processing this information, and making informed decisions. In addition to this core driving task, Mission Specialists must also make decisions to engage and disengage our self-driving system. This is why we invest in their training, monitor their performance, and provide regular feedback and coaching for continuous improvement. Mission Specialists' training is covered in more detail in [section 06.02](#).

Policies

– We have implemented a number of technologies and policies for Mission Specialists to assist with the safety of self-driving vehicle operations.

Hours of Service

We implement an Hours of Service policy informed by FMCSA Hours of Service Regulations⁵⁹ and public fatigue management research.⁶⁰ While regulations and most current research applies directly to commercial vehicle use, we believe this research is relevant to self-driving vehicle operation, in light of the complexity of self-driving systems and the attention required to maintain control of the vehicle during testing.

Our policy requires that:

- Mission Specialists confirm that they have gotten sufficient sleep to perform their duties in the course of completing their pre-mission checklist.
- Managers are trained using U.S. DOT's Drug and

⁵⁹ FMCSA, 2017, 'Summary of Hours of Service Regulations.'

⁶⁰ North American Fatigue Management Program, ND, 'North American Fatigue Management Program: A Comprehensive Approach for Managing Commercial Driver Fatigue.'

Human-Machine Interface (continued)

- Alcohol Supervisor Guidance⁶¹ to look for signs of tired or impaired Mission Specialists.
- Managers must approve any over-time hours for in-vehicle work beyond an eight-hour workday.
- Mission Specialists take a mandatory lunch break of at least 30 minutes at the midpoint of their shift and two 15-minute breaks during the middle of the second and third hour of continuous operations.
- Mission Specialists work fewer than 50 hours in a rolling seven-day period; Mission Specialists are prohibited from operating a vehicle until this metric drops below 50.
- Pilots are limited to four hours behind the wheel in a given workday and two hours without taking a break or switching positions.
- Mission Specialists rotate shifts between in-vehicle and out-of-vehicle work tasks, targeting roughly half of working time out of the vehicle.
- Mission Specialists are encouraged to alert their manager in the event they do not feel fit for planned duties.

Cell Phone Use

Mission Specialists are prohibited from interacting with their mobile devices while the vehicle is in motion or stopped in traffic. Our policy calls for a violation of this prohibition to result in discipline up to and including termination.

Monitoring

All of our self-driving vehicles are equipped with a third-party driver monitoring system. If the system detects distracted driving, an audible alert sounds in

the cabin and a notification is simultaneously sent to our remote monitoring team for review and escalation. We have also introduced an audible alert whenever the speed limit is exceeded when the vehicle is operating in manual mode.

This third-party monitoring system records acceleration, braking, cornering and tailgating events and sends this data to a specially-trained team for review. This information makes it possible to provide evidence-based feedback to Mission Specialists on their decisions.

Front Seat Touchscreen

All of our self-driving vehicles are equipped with a touchscreen tablet that communicates important information to our Mission Specialists, including turn-by-turn directions and self-driving system mode.

We follow NHTSA's Human Factors Guidance for Driver-Vehicle Interfaces⁶² to minimize distraction connected to installed vehicle components, and have established a complementary set of policies to protect against inappropriate use. Today, the touchscreen:

- Does not require input from the Pilot while driving.
- Restricts available functionality when the vehicle is traveling at speeds over 5 miles per hour.
- Minimizes use of text, background information, and options for interaction.
- Uses audio and user interface transitions and map motion to clarify information presented.

⁶¹ U.S. DOT, 2015, 'Drug and Alcohol Supervisor Training Guidance.'

⁶² NHTSA, 2016, 'Human Factors Design Guidance For Driver-Vehicle Interfaces.'

Human-Machine Interface (continued)

- Employs a visual system focused on color, iconography, and visual layout to improve glanceability.
- Optimizes color for time of day.

→ Looking Forward

Rider Experience

Rider trust is key to the successful adoption of self-driving vehicles. At this stage in our development process, our primary rider experience goal is to build and maintain trust.

As we continue to develop our rider experience, we are focused on providing:

- **Transparency**
When a rider enters the vehicle, we intend to have a touchscreen tablet in the backseat welcome them, ask them to confirm their destination, and show the vehicle's route. During the ride, we intend to enable the rider to monitor trip progress or view a visualization of the car's perceived environment on the touchscreen.
- **Control**
When requesting a ride, we intend to notify a rider that they have been matched with a self-driving vehicle and give them the ability to opt out. We intend to allow the rider to control when the trip begins and retain the option to request a stop at any time using controls on the backseat touchscreen. When a stop is requested, we intend to have the vehicle stop when and where it is safe to do so.

- **Comfort**

In addition to following the rules of the road, we intend our self-driving vehicles to provide a comfortable rate of acceleration and avoid harsh braking. We expect the backseat touchscreen interface to provide the rider with information they need, but also allow them to disengage from the touchscreen experience when desired by, e.g. dimming or restylizing the interface for nighttime rides.

- **Remote Assistance**

In the future, when our self-driving vehicles do not have Mission Specialists in the front seats, riders may need advice on accessing the vehicle, have questions about the self-driving technology, or require assistance in the event of an emergency. Additionally, they may require remote reminders about adhering to in-vehicle policies and local laws, e.g., where passengers should sit and proper restraints, like child and booster seats. We are developing remote Rider Assist functionality to facilitate some of these anticipated needs.

In the event of a crash, Mission Specialists are trained to make sure the area is safe and check on the passengers. In the future, we expect our Remote Assistance system to automatically connect to the vehicle to communicate with passengers, provide them with status updates, and facilitate further in-person assistance. We intend to have the Remote Assistant stay connected until first responders arrive and coordinate with them as needed.

- **Pedestrian and Law Enforcement Interaction**

We are developing an external speaker and microphone system to allow a Remote Assistant to communicate with law enforcement and pedestrians as needed.

Human-Machine Interface (continued)

Transitioning Between Self-Driving and Manual Driving

Transitions to and from manual mode help facilitate safe testing and help to manage a number of types of risk. These transitions are only completed by Mission Specialists who have received the training necessary to understand how and when to do so safely. Similarly, Mission Specialists must be able to easily transition out of self-driving mode and into manual mode whenever necessary to ensure safe operation.

We have designed the self-driving system to have multiple means of shifting between manual driving and self-driving modes.

- **Shifting Into Self-Driving Mode**

While the vehicle is in park at the start of a mission, Mission Specialists authenticate to Uber infrastructure and verify a two-factor challenge. This authorization and authentication is intended to ensure appropriate access prior to enabling the self-driving system. Once authenticated, the Mission Specialist manually maneuvers the vehicle into the ODD. When the vehicle is in the ODD, the Mission Specialist is informed of system readiness via visual cues on the touchscreen tablet.

In current generation vehicles, once both the system and the Mission Specialists are ready, the Pilot can shift into self-driving mode by depressing an engage button located in the center console. In next generation vehicles, the Mission Specialist in the driver's seat can shift into self-driving mode by pulling both steering wheel shift paddles simultaneously. Upon shifting into self-driving mode, audio and visual cues confirm the successful transition.

- **Shifting Out of Self-Driving Mode**

At any time, the Pilot can shift out of self-driving mode using any of the following methods:

- Depressing the accelerator pedal
- Depressing the brake pedal
- Turning the steering wheel
- Depressing the red emergency button in the center console.

While not intended as primary disengagement methods, a Pilot disconnecting their seatbelt or opening their door will also result in shifting out of self-driving mode.

- **Knowing Which Mode is Active**

In each discrete self-driving system state (manual, ready, self-driving), the current operational mode is displayed to the Mission Specialist on the front seat touchscreen using a persistent banner that changes color and text depending on the mode. Turn-by-turn instructions route lines on the touchscreen match the color of the banner. In current generation vehicles, the operational mode is also indicated by a redundant LED light strip.

- **Knowing When the Mode Changes**

As described above, when the system shifts into self-driving mode, audible and visual cues are presented to the Mission Specialist. Similarly, when the system returns control to the Pilot, audible and visual cues are presented. □

Data Recording

Our self-driving vehicles capture significant quantities of environmental and systems data every second. We use this high-resolution data in a number of ways, including system performance analysis, quality assurance, machine teaching and testing, simulated environment creation and validation, software development, human operator training and assessment, map building, and validation.

Data Types

Our self-driving vehicles record telemetry, control signals, Controller Area Network (CAN) messages, system health (e.g. hard drive speeds, internal network performance, and computer temperatures), as well as sensor and camera data.

Logging and Storage

This data is captured in real-time on the vehicle and then offloaded to our data centers for storage, cataloging, review, and labelling.

We are developing onboard data storage with reliability and resilience in mind. We verify each vehicle's data-logging capabilities and storage sufficiency before operation. Where appropriate and without risking already stored data, the onboard storage volumes perform continuous self-assessments, including monitoring read/write errors and disk fault detectors.

Data storage and processing of vehicle data will pose significant burdens to operations at scale. Nonetheless, in addition to data logged by the vehicle's Event Data Recorder (EDR), all logging modes in our system

provide for the baseline of data required for crash reconstruction, as indicated by NHTSA's guidance.⁶³

All vehicles are equipped with a backup battery to improve the system's ability to log data in the event of a crash. In the event of power failure, data is logged to solid state hard drives with the expectation that data will be written out of cache and onto non-volatile storage.

Transmission

In addition to our regular onboard storage and the EDR, vehicles transmit a small amount of data Over-the-Air (OTA) to Uber servers to provide real-time insights into how our vehicles are performing, where they are, and their current state.

Data that is transmitted OTA may also be logged to onboard storage for later cross verification and data mining. All OTA communication seeks to mitigate cybersecurity risk while communicating over redundant cellular networks provided by multiple carriers.

→ Looking Forward

Vehicle power is shut off shortly after an impact. We are developing capabilities to offload emergency data using battery backup cellular devices to transmit certain telematics and other relevant data. □

⁶³ NHTSA, 2017, 'Automated Driving Systems 2.0: A Vision for Safety.'

Crashworthiness

We evaluate crashworthiness by examining the degree to which a vehicle will protect its occupants from the effects of a crash. We are able to promote crashworthiness of our self-driving vehicles in two manners: through the crashworthiness of the base vehicle and crashworthiness of the self-driving system.

Base Vehicle

Crashworthiness of the base vehicle is defined by the vehicle structure, occupant restraint systems, and other factors.

Volvo supplies vehicles to Uber which serve as the base for our self-driving technology. To preserve the strong safety benefits of the Volvo XC90, we do not remove crashworthiness features from the vehicle. Prior to Uber's purchase, Volvo has certified these vehicles as meeting relevant Federal Motor Vehicle Safety Standards (FMVSS). These FMVSS cover everything from brake lamps to windshield wipers to the vehicle's performance in the case of front impact, side impact, and rollover.

Self-Driving System

As described in [section 05](#), we add sensors, wiring, and computers to the base vehicle to enable our self-driving system. We evaluate these modifications to avoid interference with the native Volvo safety equipment, and work closely with Volvo to avoid inconsistencies.

Modifications to the base vehicle are designed to preserve safety and structural integrity, while minimizing risk to passengers in the event of a crash.

- **Sensor Wing**
The sensor wing is mounted to the roof using modified roof rails. In the current generation vehicles, roof rails are Uber-installed; in next generation vehicles, they are factory-installed.
- **High Voltage Wiring**
Additional high voltage wiring is integrated into the base vehicle to power the self-driving computer and other devices. These cables are routed behind fixed interior panels to make them inaccessible to passengers, and include high voltage interlock protection to minimize the risk of electric shock to passengers, first responders, or others that may come into contact with the vehicle following a crash.
- **Low Voltage Wiring and Sensor Cleaning Tubing**
Low voltage wiring and other tubing has been added to the vehicle to connect the sensor wing to the self-driving computer and fluid and air compressors. Routing pathways through the vehicle have been selected in order to minimize risk of degrading airbag function.
- **Self-Driving Computer**
In current generation vehicles, the self-driving computer is housed in the rear cargo space of the vehicle. In next generation vehicles, the self-driving computer is reduced to about the size of a medium-sized suitcase and is housed beneath a tamper-resistant, load-bearing floor in the trunk space of the vehicle. □

Post-Crash Behavior

Should a crash occur, the base Volvo platform performs a variety of safety actions depending on the type of collision detected:

- **Passive Safety Features Activation**
Deploys front and side curtain airbags, activates seat belt tensioners, and automatically unlocks doors.
- **Post-Impact Braking**
Brings the vehicle to a controlled stop after the collision to avoid the vehicle entering the path of other vehicles.
- **High-Voltage Battery Disconnection**
Disconnects the high-voltage battery to minimize risk of electric shock to passengers and first responders.
- **Hazard Lights Illumination**
To warn other approaching drivers of the hazard.
- **Emergency Services Notification**
An emergency call is automatically made to personnel trained to immediately assist. The GPS location of the vehicle is automatically sent so that first responders can be called directly to the scene of the accident.

Mission Specialists' Role

As part of manual driving training, Mission Specialists undergo Incident Response Training on how to appropriately respond after an incident and engage with emergency personnel. Where it is safe to do so, the Mission Specialist remains with the vehicle post-crash to provide reasonable assistance to involved parties, law enforcement, and medical professionals. Mission Specialist training is covered in more detail in [section 06.02](#). □

Vehicle Cybersecurity

In addition to physical safety scenarios, we also consider and defend against common behaviors of actors seeking to access our systems as well as to alter and/or remove data. Self-driving vehicles interact across multiple information, network, and hardware domains, thereby giving rise to a number of possible threats from malicious actors.

Uber's current fleet of self-driving vehicles is built on base vehicles designed for human drivers and therefore may contain component-limitations and communication-designs limiting active security measures. Because of these potential threats, a robust cybersecurity program is critical to promote the safe deployment of self-driving vehicles on public roadways. We have incorporated security mechanisms into the self-driving computer, sensor components, our software, and interactions with the base vehicle to reduce daily operating risks as well as in the event of attempted action by an unauthorized party. These security controls are integrated with individual components and incorporated within the platform design to defend against potential threats. For example, authorized entities must be authenticated before they can deploy firmware and software, interact with APIs, or access metadata.

Our cybersecurity approach is informed by best practices described by NHTSA and relevant industry groups, including ISO, SAE International,⁶⁴ and the Automotive Information Sharing and Analysis Center (Auto-ISAC).⁶⁵ As recommended by NHTSA, Uber adopts and designs controls with the expectation that high-risk domains (e.g. cellular-adjacent devices) may be occupied or manipulable by malicious actors.

Uber has designed and is employing security-specific principles, controls, and technologies within the self-driving computer, vehicle platform, and network infrastructure as detailed below.

⁶⁴ SAE, 2012, 'Cybersecurity Guidebook for Cyber-Physical Vehicle Systems J3061.'

⁶⁵ Auto-ISAC, 2015, 'Best practices.'

Vehicle Cybersecurity (continued)

Hardware Security Controls

Key Management

The Uber self-driving vehicle utilizes asymmetric cryptographic primitives to establish trust with remote entities and enable trusted execution. In order to securely manage these cryptographic primitives, we have incorporated hardware security modules within the various security domains on the vehicle platform. These hardware devices integrate with core system components and services to provide protections such as remote attestation, domain-domain authentication, and a secure mechanism to validate executable firmware, code, and data originating from other security domains.

Functional Separation

The vehicle, sensors, and compute platform are designed to have distinct communication domains which separate a sensor, device, computer, or remote service operating within the larger system. Self-driving vehicles are comprised of many devices; each of the components onboard and the network edge systems are within their own security domain. Strongly-controlled security domains help, for example, to isolate sensor components from vehicle command, and rider experience devices from motion-control devices. Communications between strongly-controlled security domains are API-driven and can feature cryptographic security restrictions (e.g. signature verification), environmental and code-execution restrictions, and authentication requirements. Some security domains on the vehicle are protected through physical means.

Secure Networking Devices

Onboard, in order to manage the risk of cellular attacks, we have implemented redundant hardware and engaged multiple internet providers and cellular

networks; these redundant domains mitigate impact in the event that a single cellular network or modem is attacked or compromised. Offboard, the vehicle's cellular modems and the Uber network terminus systems each have network-layer and hardware security devices and controls to help isolate these high-risk components from other on-vehicle and datacenter components. Vehicle messaging is controlled with cryptographic, logical, and policy-based approaches.

Security Architecture

Cryptographic Signatures

Autonomy pipelines and many of the motion control devices on Uber's self-driving platform run on execution environments where code is executed only after we have authenticated the code's origin as a trusted source using cryptographic sources of identity. Cryptographic signatures also enable traceability for artifacts of interest by tying artifacts to specific users and with the processes and systems which generate firmware, software, maps, and models. These cryptographic signatures also enable processes to attest to validity and correctness, which is critical for artifact association (release processes), sensor calibrations, and configurations. Consistent with cybersecurity best practice,^{66, 67} these cryptographic signatures also provide a valuable integrity check by enabling code/firmware verification during updates and prior to execution. This methodology protects devices from malicious modification and helps to mitigate the risk of an adversary modifying the firmware to persistently perform activities on their behalf.

Mission Specialist and Data Access Control

Daily operations involving Mission Specialists on

⁶⁶ U.S. Department of Commerce National Institute of Standards and Technology (NIST), 2018, 'Platform Firmware Resiliency Guidelines.'

⁶⁷ U.S. NIST, 2011, 'BIOS Protection Guidelines.'

Vehicle Cybersecurity (continued)

the track and on the road are strictly controlled through security workflows which require multi-factor authentication and functional authorization for specific missions. In certain elevated privilege circumstances, short-lived certificates will be issued after multi-factor user authentication has completed for the specific vehicle that needs to be accessed. This more privileged access generates logs which detail actions taken by the authenticated user or tool. In normal operation, our self-driving vehicles transition frequently between operational and functional modes, e.g. from manual to self-driving mode, from not serving to serving riders. The previously described authentication and monitoring mechanisms are used to mitigate risks during these mode changes.

Onboard Communication Security

The self-driving computer employs best practice⁶⁸ transport layer security when communicating over internal networks. This ensures that critical autonomous decisions are made over integrity- and intercept-protected channels. For example, should an adversary obtain access to the self-driving computer networks through a sensor, the self-driving computer would not be subject to command forgery.

Remote Network Access Policies

A self-driving vehicle needs to be able to communicate with the data center through a secured network. These communications must be resilient to a broad spectrum of attacks relevant to any mobile network. The communication layers employ best practice secure protocols to protect the channel and transmitted data from interception and modification. We use secure tunnels gated with regularly audited short-lived vehicle-specific certificates for system telematics and communication.

Secure Software Engineering

– Minimizing Attack Surface

Our vehicle security development policies focus on attack surface reduction at every cross-domain interaction layer by requiring specific protocols and published API definitions. Protocol-level controls inform the additional security constructs required within each domain and between domains. Security domains that include inherently risky protocols can be augmented with security devices to manage risk. When possible, the principle of ‘least privilege,’ meaning access only as required for authorized purpose, is actively applied to minimize the risk presented by a single weak component and its immediate neighboring domains and components.

Adversarial Simulation

The vehicle security team collaborates internally and with our partners to identify, document, and remediate weaknesses in hardware, software, protocols, APIs, and overall platform risks. These simulations and reviews are designed to evaluate the interactions of components at a platform level, identify any weaknesses associated with their incorporation, and evaluate platform security features and components. We document and evaluate risks to the vehicle platform and self-driving computer in order to recommend security improvements to individual components and improve platform-level security controls.

→ Looking Forward

Like others in the industry, we continue to explore and invest in developing improvements for the security mechanisms, policies, and components described above. In collaboration with our partners, we are incorporating security improvements across the

⁶⁸ U.S. NIST, 2017, ‘Guidelines for the Selection, Configuration, and Use of Transport Layer Security Implementations.’

Vehicle Cybersecurity (continued)

vehicle platform and are dedicated to pushing new security features into components to improve the security posture for future self-driving vehicles across the industry. Below are a few of the areas where we are prioritizing our research.

- **Software and Firmware Signatures**
We are working to ensure future vehicle platforms adhere to strict update and runtime signature verification requirements of firmware already implemented by the self-driving computer today.
- **Hardware Isolation of Functional Domains**
We are developing devices that we hope will further improve device and component isolation. These components are designed to be hardware risk reduction mechanisms to help protect disparate security zones and to provide a single-point of application-level ingress/egress for any interaction.
- **Functional API Separation**
We are continuing to establish clearly defined APIs that traverse functional domains and reduce the possible message permutations that require testing, thereby enabling more complete evaluation of protocol misuse across domain boundaries both during runtime and during development when protocols change. □

Uber's Self-Driving Safety Principles

06.05 → Trustworthy

Principle 5 Trustworthy



We recognize the importance of earning the trust and confidence of both the public and various levels of government in support of successful development and deployment of self-driving vehicles. We are committed to earning and maintaining that trust with our stakeholders - riders, government officials and policymakers, non-governmental advocacy and interest groups, industry, partners, employees, and the general public, which includes riders, drivers, and couriers on the Uber platform.

We believe that the most effective approach to building trust is to provide regular, consistent, accessible information on our development efforts, business plans, and the potential impacts of our

operations on local communities where we operate. We want our stakeholders to have high-quality information on the technology in order to make informed decisions about use and regulation of self-driving vehicles. The detail provided through this report represents one mechanism to supply such information.

Trustworthy covers the following NHTSA safety elements: *System Safety, Consumer Education and Training, and Federal, State, and Local Laws.* □

Stakeholders shall be engaged participatively and provided appropriate verifiable or audited evidence of safety.

- Our *stakeholders* include riders, regulators, and legislators, along with all people with whom we share public roads and organizations that advocate on their behalf.
- We recognize that taking steps to inform the public about our approach to safety and how we are working on self-driving vehicles is necessary and important for building trust. *Participative engagement* means consulting and partnering with stakeholders to understand their needs, and to revise our approaches to best reflect this broader set of interests.
- We understand that we cannot simply provide descriptions of the safety performance of our systems. We are committed to employing various methods to *provide evidence of safety performance.*

Uber's Self-Driving Safety Principles

06.05 → Trustworthy

System Safety

Safety Performance Metrics

We have initiated an effort to develop a set of valuable performance metrics for communicating self-driving vehicle safety effectively, both internally and externally. A number of metrics, such as disengagements and self-driving miles traveled, have been taken up by self-driving pilot program designers and self-driving developers as indicators of progress in development. We are concerned that overemphasis on these metrics may create perverse incentives to, e.g. avoid disengagement even in scenarios where it is the safe choice and/or be applied and defined inconsistently across developers. These metrics, if applied consistently within the efforts of an individual developer, may provide some useful information on improvement over time, but they should not be considered appropriate objective, cross-developer safety metrics in deployment.

We are undertaking internal development work and collaborating with researchers and partners to establish a set of safety performance metrics which:

- Are specific to the stage of development (development, testing, deployment).
- Are specific to one or a number of agreed ODDs, or a standard set of capabilities or scenarios.
- Have comparator metrics for human-driven vehicles.

Independent Experts

We believe that engaging independent experts to review our safety approaches and performance is essential to our learning and development. These third-

party reviews can also provide additional confidence to our customers, government officials, and others while self-driving technology is in development.

We have already undertaken a number of external reviews, including two 2016 assessments by an independent assessor of the appropriateness and effectiveness of our safety measures and alignment with ISO 26262, and a 2018 review of ATG's safety culture by a team of external experts, including a former chairman of the National Transportation Safety Board. We intend to prioritize these kinds of reviews, actively consider their recommendations for incorporation into our programs, and look to share the results when appropriate.

We expect these independent reviews to consider particular elements of our safety approach, rather than assess our entire system for safe performance. This approach allows us to prioritize review efforts and engage experts with specific expertise and competence in particular areas.

→ Looking Forward

Safety Standards

As described in [section 04](#), Uber is preparing to bring self-driving technology to the world by hosting other developers' self-driving vehicles on the Uber network. In order to do so confidently, we are developing a standard qualification process to be used across all parties that look to deploy self-driving vehicles on the Uber network. Such safety standards may include, e.g. measures of system safety in relevant environments, a process for identifying and describing ODD and ODD-relevant capabilities, and performance requirements in ODD-relevant scenarios. □

Uber's Self-Driving Safety Principles

06.05 → Trustworthy

Consumer Education and Training

Safety Reports

- We believe that the voluntary safety reporting envisaged by NHTSA's 2017 autonomous vehicles guidance is an important platform for self-driving technology developers to communicate consistently and regularly regarding progress in development, remaining challenges, and plans for deployment.

This safety report will be the first in a series of regular updates, released at key points of transition and development of our self-driving system.

Public Engagement

- We will proactively seek to:

- Educate consumers on safety features and interaction with self-driving vehicles through a number of channels, including blog posts, marketing campaigns, and direct exposure to our self-driving vehicles. We intend to use these channels to explain the technology underlying our self-driving system and relay first-hand accounts of the rider experience.
- Inform and engage with the communities where we operate by organizing community events, holding town hall sessions, providing notices of operation, and collecting feedback.

Stand-Ups for Safety

- Putting safety first means ensuring that every member of our team understands how we define and

approach safe development and performance, and can bring this understanding to bear in their work, whether as a Mission Specialist, software engineer, or communications or marketing professional.

We want to create the space and time for this active engagement with all staff on safety. We plan to hold regular Stand-Ups for Safety, whole- or half-day events during which normal work will be suspended and all staff should have the opportunity to learn about our Safety Principles and approaches, and from external safety experts on developments in the field. Over time, we plan to invite other developers to run similar programs concurrently, and invite members of the public, government officials, and others to participate.

→ Looking Forward

Self-Driving Safety and Responsibility Advisory Board

- Uber ATG is working to establish a self-driving safety and responsibility advisory board, comprised of independent experts, to provide objective reviews of and input onto aspects of our self-driving program. We hope that the board will consult, advise, and review Uber's approaches at the intersection of self-driving, mobility, safety, and company responsibility in regards to both development and production-level self-driving systems.

We plan for the board to consist of a panel of independent, external experts with objectives to:

- Review and advise on Uber's self-driving program's policies, culture, operational procedures, and processes.

Consumer Education and Training (continued)

- Review Uber's approaches on broader industry topics, such as readiness for fully driverless operation, and responsibility on issues such as technology ethics and sustainability.
- Identify potential risks and hazards and recommend corrective actions. □

Federal, State, and Local Laws

Federal Motor Vehicle Safety Standards

—

Prior to Uber's receipt, Volvo has certified base vehicles as meeting all applicable FMVSS.

We recognize that new FMVSS and/or changes to existing FMVSS may be promulgated in coming years to take account of unique features of self-driving vehicle design and capability. We welcome these further clarifications and hope to engage collaboratively with the U.S. DOT and NHTSA as they seek to prioritize, develop, and implement these standards.

State and Local Laws

—

Rules of the road are usually set by states and localities. Safe deployment of self-driving vehicles requires attention to these rules to facilitate the integration of self-driving vehicles with the broader set of actors in the transportation environment.

For specific capabilities we design for, we assess the relevant traffic laws for a given ODD to ensure that those rules are integrated into the self-driving system. As with other limitations on the behavior of the self-driving system (see section 06.01), we enact formal limits within the self-driving system to promote compliance with these rules. For example, we build state and local road rules, e.g. speed limits, into our high-definition maps and program the vehicles to follow these rules, e.g. by staying below the prevailing speed limit. In addition, we intend to insure all vehicles in accordance with the insurance and financial responsibility laws of the states where they are operated. □

Disclaimer

This report, including but not limited to information contained in the sections labelled *Looking Forward*, contains management's current intentions and expectations for the future, all of which are forward-looking statements. The words "estimate," "plan," "may," "intend," "expect," "believe," "anticipate," and similar expressions are intended to identify forward-looking statements. Actual results may differ materially from these forward-looking statements due to various factors. There can be no guarantees that forward-looking statements will be true. You should not place undue reliance on forward-looking statements, which speak only as of the date of this release. □

APPX →

List of Acronyms

Anti-lock Braking System (ABS)
 Application Programming Interface (API)
 Automated Driving System (ADS)
 Automatic Emergency Braking (AEB)
 Automotive Information Sharing And Analysis Center (Auto-ISAC)
 Automotive Open System Architecture (AUTOSAR)
 Battery Electric Vehicle (BEV)
 Controller Area Network (CAN)
 Dynamic Driving Task (DDT)
 Electronic Stability Control (ESC)
 European Commission (EC)
 The European New Car Assessment Programme (Euro NCAP)
 Event Data Recorder (EDR)
 Fault Injection Training (FIT)
 Federal Motor Vehicle Safety Standards (FMVSS)
 Global Positioning System (GPS)
 Hardware In the Loop (HIL)
 Inertial Measurement Units (IMU)
 Insurance Institute for Highway Safety (IIHS)
 Institute of Electrical and Electronics Engineers (IEEE)
 International Energy Agency (IEA)
 International Organization for Standardization (ISO)
 International Organization of Motor Vehicle Manufacturers (OICA)
 Light Detection And Ranging (LIDAR)
 Motor Industry Software Reliability Association (MISRA)
 National Aeronautics and Space Administration (NASA)
 National Safety Council (NSC)
 Object and Event Detection and Response (OEDR)
 Object Detection and Classification (ODTAC)
 Operational Design Domain (ODD)
 Operator Control Station (OCS)

Original Equipment Manufacturer (OEM)
 Over-the-Air (OTA)
 Plug-in Hybrid Electric Vehicle (PHEV)
 Track Validation Testing (TVT)
 U.S. Department of Defense (DOD)
 U.S. Department of Transportation (DOT)
 U.S. Environmental Protection Agency (EPA)
 U.S. Federal Motor Carrier Safety Administration (FMCSA)
 U.S. National Highway Traffic Safety Administration (NHTSA)
 U.S. National Institute of Standards and Technology (NIST)
 U.S. National Transportation Safety Board (NTSB)
 Uber Advanced Technologies Group (Uber ATG)
 Ultrasonic Sensors (USS)
 Vehicle Interface Module (VIM)
 World Health Organization (WHO)