Starsky Robotics – Voluntary Safety Self-Assessment

The National Highway Traffic Safety Administration (NHTSA) has defined a set of safety design elements describing priority topics that any manufacturer or designer of an automated driving system (ADS) should consider when developing an automated vehicle (AV).[1] NHTSA encourages industry to provide the public with information demonstrating their incorporation of these elements in the form of Voluntary Safety Self-Assessments (VSSAs) as part of a best practices approach to introducing ADSs to U.S. roadways. Starsky Robotics welcomes the opportunity to publish our VSSA, which describes our commitment to safety in developing ADS-equipped commercial motor vehicles and our approach to the design elements outlined by NHTSA.

Starsky is designing a deterministic automation system that utilizes a human-in-the-loop for certain decision-making processes and completes off-highway segments of long-haul trucking routes by exercising direct remote control over commercial motor vehicles (CMVs) via telemetry. This unique, complimentary combination of human decision-making and automation allows Starsky to execute end-to-end trips and efficiently haul freight without a human driver physically present in the cab of the truck. Importantly, this VSSA is an iterative document that characterizes the current capabilities of our system. As our system evolves, we will update our VSSA accordingly. Starsky believes that safety is the most important design goal for a successful ADS. This VSSA is one part of our implementation of a proactive approach to working with the public and government stakeholders to bring safe automated commercial motor vehicles (ACMVs) to America's roads.

I. Introduction

Starsky Robotics is a San Francisco-based startup developing ACMVs that are autonomous on the highway and remote controlled, or teleoperated, by human drivers for the first and last mile from highway exit to distribution center. We received our operating authority in 2017 and regularly haul freight to test our systems using weighted loads. In February 2018, we became the first company to publicly test an ACMV, monitored by a remote teleoperator, without a driver physically present in the cab on a closed road in Florida.

The fundamental thesis behind Starsky is that the long-haul truck driver shortage is real. According to the American Trucking Associations (ATA), the U.S. trucking industry is currently experiencing a shortfall of approximately 50,000 commercial motor vehicle (CMV) drivers, a shortage

[1] https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/13069a-ads2.0_090617_v9a_tag.pdf

that is expected to increase to 174,000 drivers by 2026.[2] ATA estimates the trucking industry will need to hire 90,000 new drivers annually to meet anticipated demand. It is increasingly difficult to find drivers who are willing to spend months on the road away from their families, particularly among younger generations. In addition, annualized driver turnover at large carrier fleets has reached and in some cases exceeded 100 percent.

The driver shortage has a significant impact on rising freight transportation costs for U.S. goods. In 2018, several major food manufacturers cited the driver shortage as a direct cause of increased costs for consumers. As 71 percent of U.S. goods are transported via CMV, representing approximately 10 million tons of freight annually, addressing the driver shortage and improving freight transportation efficiency is vital to the future of the U.S. economy.

At Starsky, we are seeking to solve the driver shortage by allowing human drivers to work in office environments while making trucks autonomous on the highway. Starsky uses teleoperation, or direct remote control of a CMV, to complete off-highway first and last mile operations and safely navigate complex, context-heavy traffic environments, such as truck yards. As remote teleoperators, Starsky's drivers are provided meaningful employment opportunities where they can utilize years of experience in the long-haul trucking industry while remaining close to their homes and families. In addition to enhanced working conditions for remote drivers, our use-case has the potential to relieve downward pressure on wages, increase safe driving practices by providing regular meal and restroom breaks, reduce annualized driver turnover, decrease incentives for drivers to push the limits of their hours-of-service requirements, and improve a rampant trend in the trucking industry whereby employees are misclassified as independent contractors.

Most importantly, exit-to-exit highway automation combined with improved conditions and safety training for teleoperators holds immense promise for the future of highway safety. According to recently released NHTSA data, crashes involving large trucks claimed 4,761 lives in 2017, a nine percent increase from 2016 levels. Our approach to ACMV deployment has the potential to reverse this trend by improving safe driving practices through well-trained, well-rested teleoperators and by using exit-to-exit highway automation to significantly reduce CMV accidents that are attributed to human error, including drunk and distracted driving.

As the rate of CMV-related fatalities tragically continues to rise, Starsky looks forward to working with government partners to fully-realize the future employment, economic, and safety benefits of the emerging ACMV industry by bringing driverless trucks to market.

[2] ATA 2017 Driver Shortage Analysis. https://www.trucking.org/article/New%20Report%20Says-National-Shortage-of-Truck-Drivers-to-Reach-50,000-This-Year

II. System Safety

Safety must be a deliberate design goal from the beginning of ADS development. Safety will not happen by accident and cannot be "added on" at the end of the design process. Starsky believes that safety can co-exist with an innovative, cutting-edge product when it is made an explicit design goal. To that end, Starsky has adopted applicable best practices, standards, and processes from other well-established industries with a history of producing safe and reliable products, including aviation, automobiles, medical devices, and the military. Each of these industries has created methods tailored to their specific risk profiles. None of these methods fit our use-case precisely, so Starsky has utilized these standards where appropriate and created a systems engineering process that fits our needs and application risk.

Safety starts with a system definition. Starsky begins by setting a goal to meet a specific use case tied to a specific operational design domain (ODD). The combination of use case and ODD allows us to create a system-level design document that outlines requirements for the Starsky system. These requirements are focused on what is needed to create the product, or feature requirements.

System safety must include risk assessment.[3] Our goal is to identify as many risks of failure as possible at the beginning of the design cycle. An identified risk can be mitigated by design features, redundancies, or narrowing our ODD. Starsky uses a process of failure mode effects analysis (FMEA) similar to what is defined in ISO 26262. The analysis is completed at the top system level, and then at subsystem levels. FMEAs are created for both software and hardware systems. The FMEAs are reviewed, critiqued, and updated as our system design evolves. Each failure mode is scored for severity, exposure, and controllability. This allows Starsky to rate the risk of each failure mode in a quantifiable manner and prioritize risks for examination and mitigation.

The FMEAs are used to drive functional safety requirements. Feature requirements describe what we want to design to make the product work properly, and functional safety requirements describe what we want to design to respond in the event that an intended function fails. Functional safety requirements may include the detection of specific failure modes, redundancy in critical systems, and fallback modes or systems for failures. Functional safety requirements may be imposed at the system level, or at a subsystem level. Functional safety engineering explicitly recognizes that no system is perfect, and that one should consider failure with deliberation and intent to control the failure.

Starsky combines Silicon Valley's innovation culture with a focus on safety. Starsky practices a fast-paced, continuous integration software development process to achieve quick-turnaround without compromising quality by incrementally checking software updates. Each request to update a portion of the software is reviewed and checked with a unit test, and a new software change must pass integration tests that check the functionality of the system. If these tests are passed, then the software

[3] Appendix contains a discussion of risk and safety.

can be deployed to software-in-the-loop and hardware-in-the-loop systems that can test the entire system performance with either simulated data or actual recorded data. This process means that Starsky can upload new software onto a research truck very quickly while maintaining confidence in the performance of the system.

A recent report by the National Academies of Sciences, Engineering, and Medicine[4] describes the importance of developing a safety culture to reduce harm from vehicular traffic. Broadly speaking, the report categorizes contributing factors of traffic accident risk into two categories: systems hazards (such as road geometry, sightlines, or equipment) and behavioral hazards (such as speeding, drunk driving, driving without seatbelts, or "clipping" red lights). Starsky can mitigate systems hazards through our ODD limitation process, described below. Our system will understand speed limits and curves and will only be deployed on roads where we have proven the system can reliably navigate the route.

Starsky can have a significant impact on behavioral hazards by deliberately adopting a conservative, "safest driver on the road" approach to our system. In 2006, the Federal Motor Carrier Safety Administration (FMCSA) released the *Large Truck Crash Causation Study*.[5] The report was based on an analysis of more than 1,000 accidents between cars and trucks, and a subset of the report describes factors that contribute to cars and trucks colliding on U.S. roadways. In over half of these accidents, trucks were coded with a "critical reason"[6], 87% of which were attributed to driver-related factors such as decision making and inattentiveness. These factors are directly under Starsky's control when designing our system. Where a human might be tempted to increase speed to hit a distance target before exceeding their allowed time on the road, a Starsky truck can make a different choice. An ACMV can be designed so that it will never exceed the posted speed limit and always follow traffic at a safe distance, rather than aggressively tailgate.

Human factors such as fatigue, distraction, or needing restroom breaks affect safety, and external factors such as the availability of parking can exacerbate these hazards.[7] Even when drivers themselves determine they should not be driving, they are often pressured to make difficult decisions because of the practicalities of driving and the demands of the trucking industry. Starsky's use of automation and remote teleoperators means that a Starsky-equipped truck will never have to decide whether to operate in the face of human pressures that could decrease safe driving practices. Rather than speeding or pushing hours-of-service limitations, in any instance where a long drive or delay would impact safety, teleoperators will simply change shifts or the truck will find a safe place to park, allowing Starsky to significantly enhance safe driving practices and achieve the operational benefits of

[4] National Academies of Sciences, Engineering, and Medicine. 2018. *A Strategic Approach To Transforming Traffic Safety Culture to Reduce Deaths and Injuries*. Washington, D.C.: The National Academies Press. https://doi.org/10.17226/25286.
[5] https://ai.fmcsa.dot.gov/ltccs/default.asp
[6] Critical reason: immediate reason for the crash; coded to one vehicle only. Crash associated factors that might be pertinent that were present during the crash.
[7] A 2016 ATRI survey on parking highlights the choice truck drivers are often forced to make when running against hours-of-service limitations and find no parking available. This happens most often when drivers have been driving all day and would be the most tired. http://atri-online.org/wp-content/uploads/2016/12/ATRI-Truck-Parking-Case-Study-Insights-12-2016.pdf

team drivers at all times. When weather, traffic, or other conditions make driving unsafe, the Starsky system can simply pull of the highway to park in a safe spot and wait for conditions to change.

III. Operational Design Domain

The ODD[8] describes the environment and conditions where the Starsky system has been designed to operate. It is the responsibility of the designer to understand what impact the environment has on the AV, including types of traffic, roads, lighting, and weather.

Starsky has chosen to start with an extremely limited ODD. The Starsky strategy is to solve simple problems and deploy systems where the ODD matches the solution. The Starsky system is not designed to operate everywhere, on all kinds of roads, in all kinds of traffic, and will only be deployed on routes that have been explicitly confirmed to be within our ODD. The Starsky ODD is generally limited to freeways[9] (limited access, without grade level crossings or cross traffic, divided or with a wide median) and short, off-highway sections at the beginning and end of a given route. Interstate highways are exemplar freeways, designed for mobility and long-distance travel. For the Starsky system to operate properly, the freeway must have lane markings that can be recognized and tracked by the system. In practice, this means the actual ODD for the Starsky system can be defined as a collection of pre-surveyed and qualified "white-listed" routes.

Furthermore, the Starsky ODD includes appropriate lighting and weather conditions. The system must be able to see and track lane lines. When weather or specific lighting conditions prevent this, the system will not drive. Currently, the system will not drive in fog or rain that is heavy enough to make lane detection unreliable.

The Starsky ODD includes the consideration of speed. The Starsky system operates as a finite state machine[10], with a collection of well-defined, deterministic allowable behaviors. Some of these behaviors are affected by the speed of the truck, and speed is therefore included as a consideration when behaviors are selected. Because Starsky relies on telemetry to and from the truck to perform the teleoperation task, described below, the ODD definition includes the consideration of LTE network connections and GPS reception. For example, a long tunnel where LTE network service is not available would be outside of the Starsky ODD.

The ODD is driving-behavior specific. Starsky's behaviors are a set of distinct, deterministic commands available to a decision-maker. This means that some behaviors may be allowed under certain specific conditions but not under others. The beginning of a route might be run using the

[8] Operational design domain is defined as part of SAE J3016 "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles."
[9] https://www.fhwa.dot.gov/planning/processes/statewide/related/highway_fucntional_classifications/section03.cfm
[10] A finite state machine is an algorithmic model designed as a set of allowed states with defined transitions between the states.

teleoperation behavior, which allows off-highway operation by having a human remotely drive the truck. The longest part of the drive on an interstate would be handled by an autonomous behavior. The autonomous behavior would be inappropriate for off-highway driving because of the more complicated traffic environment. Teleoperation would be inappropriate for interstate driving because the autonomous system is better at executing mechanical tasks like lane keeping and adaptive cruise control (ACC) in real-time.

Starsky treats our ODD as an operational constraint. When the Starsky system detects a violation of our ODD, it will achieve a minimal risk condition (MRC). Our ODD can be broadened as we demonstrate additional capabilities or as we whitelist additional routes. However, the concept of having a very limited ODD is central to our risk planning, and as the ODD changes, Starsky will change our risk analysis accordingly.

IV. Object and Event Detection and Response

Object and Event Detection and Response (OEDR) describes what SAE J3016 refers to as the tactical-level decisions made when driving.[11] OEDR captures the requirement to recognize surrounding traffic, monitor relative positions and speed, and generally perceive the driving environment. OEDR then drives the operational level of control, executing a specific trajectory.

For all AVs, this is arguably the most complicated part of the automation system. Humans routinely navigate city traffic that includes a mix of cars, motorcycles, bicycles, trucks, pedestrians, and scooters; complicated traffic controls that include signs and signals, that may be permanent or temporary; and traffic flowing with the vehicle and at angles to the vehicle, all in environments where humans are liable to violate traffic rules. This complicated real-world traffic environment is the reason that AVs designed for city use are still in development after years of testing and billions of dollars spent. Starsky is taking a starkly different approach to solving the OEDR problem.

Starsky's specifically limited ODD informs our OEDR strategy. When driving in relatively unstructured environments at slow speeds, such as truck yards, a teleoperator will be responsible for the OEDR task just as if they were driving the truck directly. The driver sits in a control station that gives the driver video feed from multiple cameras. The driver can see ahead and to the sides as if they were physically present in the cab of the truck. Other cameras give views rearward to remove a blind spot that is typically present alongside the truck in legacy CMVs. Routes are designed so that the difficulty of the teleoperation task is minimized, including through the consideration of network quality, unprotected left-turns, typical traffic, and speed along the route.

For high-speed driving, the Starsky system uses a combination of camera images and radar returns to detect objects and vehicles in the environment around the truck. Camera images are

[11] SAE J3016 is the widely referenced industry standard that defines terms related to automated driving, including the Level 0 – Level 5 classification of automation capabilities.

processed at high frame rate through a machine-learning network that has been trained on highway traffic. This combination of vision and radar allows the system to identify the location and speed of the surrounding vehicles. The system recognizes traffic, objects in the lanes, and objects on the shoulder and will automatically control its trajectory and speed to avoid collisions. The remote driver will monitor the system and is responsible for helping navigate complex, context-driven environments.

V. Fallback (Minimal Risk Condition)

Starsky has designed a set of emergency behaviors that can be executed to place the truck into a MRC. We have a variety of MRCs available and the system will prioritize the safest possible fallback. Starsky has implemented a sophisticated set of system diagnostics that runs in real-time whenever the system is on. Starsky uses this set of diagnostics to judge at every moment which fallback behaviors are possible and best-suited for execution. This allows an immediate response to any system failure, the detection of an anomalous situation, or a violation of the Starsky system ODD. The system will determine which emergency behavior is most appropriate for a given driving scenario within a designated number of milliseconds. These emergency behaviors are designed for the first generation of our system, which includes a teleoperator monitoring a single truck from a remote station for the entirety of a given route. As our system evolves, the process for achieving an MRC will change and will be reflected in our VSSA accordingly.

Currently, the first priority MRC slows the truck and places the system under direct teleoperation control. Teleoperation requires the truck to have a good video and data connection to a remote-control station. The truck must also be able to continue to execute lane keeping until the truck slows enough for the teleoperator to take over safely. If that is not possible, the system will execute the most appropriate maneuver that brings the truck to a stop in a safe location. Under the second priority MRC, the system will achieve maximum deceleration, pull the truck onto a shoulder, and gradually come to a stop. Pulling off on the shoulder requires working cameras and radar on the side of the truck that is closest to the shoulder. The system will check for the presence of a shoulder, as well as for objects on the shoulder as it slows and pulls off the road. Allowing a teleoperator to safely get the truck to the nearest exit or pulling the truck to the shoulder of the road is the preferred MRC. In advance of any deployment, checking for adequate shoulder availability is part of our selection process in whitelisting routes. In the unlikely event that the truck is unable to execute these preferred fallbacks, Starsky has also designed worst-case scenario MRCs to bring the truck to a stop as quickly as possible. In such an instance, "controlled stop" means that the system has the capability to control both steering and braking, so that the system can continue to ensure it is not leaving the lane as it brakes and is not hitting objects in the lane ahead. "Immediate stop" is the final fallback mechanism for the extremely improbable occurrence that everything in the system is failing.

Table 1. Fallback Behaviors

Priority	Behavior	Comments
1	Slow down and handoff to teleoperator	First priority fallback; allows a teleoperator that has been monitoring the truck to decide whether to continue the drive or to pull over in a safe spot.
2	Slow down, pull onto the shoulder and stop	Fallback if teleoperation is not functional, driver does not take control, or driving scenario requires immediate response.
3	Controlled stop	Fallback if a clear shoulder is not found (shoulder is not present, object detected in shoulder, or side sensors failing).
4	Immediate stop	Fallback if whole system fails; a solenoid will trigger the air brakes.

For upcoming deployments, Starsky will have a nearby responsible employee who will reach and move a stopped truck as quickly as possible. We work with government and law enforcement in advance of any deployment to take steps to monitor traffic conditions, prepare to warn nearby traffic, and otherwise mitigate the risk of a statistically unlikely, worst-case scenario event. In all cases, the system maintains the ability to control turn and hazard signals. If the truck needs to execute a fallback behavior, the system will use appropriate turn signals or hazards.

The Starsky system will operate without a driver physically present in the cab of the truck. This allows the Starsky system to implement fallbacks that are impractical or simply impossible for a manned vehicle. For example, the Starsky system can simply pull over and wait if conditions deteriorate in a manner that make the environment unsafe to continue driving. Imagine if an automated taxi wanted to pull off the highway for two hours because it encountered unforeseen adverse weather conditions – this is not a practical behavior for a vehicle carrying passengers.

In addition, when presented with an otherwise unavoidable collision, we can deliberately execute a maneuver that avoids nearby traffic and people without concern for the safety of a human occupant in the truck. This allows Starsky to prioritize the safety of surrounding traffic over the unmanned truck. This choice would not be possible for a manned system or an ADS carrying passengers. When designing fallbacks for achieving an MRC, Starsky needs to consider the possible risk the truck poses to others on the road and in the surrounding area, but does not need to consider the risk to any people in the truck itself.

VI. Validation Methods

The RAND Corporation has published very thoughtful studies about the validation of AVs. The RAND report *Driving to Safety*[12] argues that AVs would have to drive an impractically high number of miles to statistically demonstrate that they are as safe as human drivers. This study does not show that AVs should be driven millions of miles before they are trusted or considered safe, but rather, that we must find a better way to validate AV safety. Requiring this immense amount of driving before allowing AVs on the road sets an unachievable validation threshold for deployment. For Starsky, our unique automation system characterized by deliberately limited ODDs, a deterministic approach to OEDR, and the use of a human-in-the-loop for certain decision-making processes means we can adopt distinctly different validation metrics.

The RAND report *Measuring Automated Vehicle Safety: Forging a Framework*[13] published in 2018 considers AV safety validation with the aim of proposing a framework of methods that might be used to build safety confidence in an AV system. The report notes that the validation data available (standards/processes/design, leading and lagging measures) changes with the project stages (simulation vs. closed road prototype testing vs. public road product testing). As a result, it is difficult to establish a single metric for validating AV safety, and companies and policymakers should instead be considering a matrix of data that will change as the AV project advances.

The standards/processes/design category captures existing industry practices and standards that might be applicable to an AV development effort. For example, Starsky follows a systems engineering process and borrows elements from ISO 26262 to ensure that safety is a concern that underlies the entire design process. Starsky can look to existing performance standards for ACC and lane keeping products to guide testing. Leading measures refer to events or actions that do not result in an accident but are indicative of conditions that could be precursors to accidents. A commonly used leading measure is disengagements[14], events when a safety driver takes control of the AV, whether because the safety driver or the AV itself judges that the system is failing to control the vehicle safely. Lagging measures include statistics on actual accidents and harm. These measures are collected for vehicles as "fatalities per 100 million miles" and set the baseline for the de facto level of acceptable risk for AVs: that they should be "as safe or safer than a human driver."[15] ACMVs should be safer than trucks driven by humans and therefore cause fewer on-road fatalities, or at very least meet an equivalent safety standard for highway accidents.

---

[12] *Driving to Safety,* Kalra, N. and Paddock, S.M., RAND Corporation, 2017. https://www.rand.org/content/dam/rand/pubs/research_reports/RR1400/RR1478/RAND_RR1478.pdf

[13] *Measuring Automated Vehicle Safety: Forging a Framework,* Fraade-Blanar, L., et al, Rand Corporation, 2018. https://www.rand.org/pubs/research_reports/RR2662.html

[14] The California DMV requires companies testing automated vehicles in the state to annually report disengagements as one example. The data is made public: https://www.dmv.ca.gov/portal/dmv/detail/vr/autonomous/disengagement_report_2017

[15] FMCSA publishes "Crash Facts" for each year. In 2016 large trucks were involved in crashes leading to 1.3 fatalities per 100 million miles driven. https://www.fmcsa.dot.gov/sites/fmcsa.dot.gov/files/docs/safety/data-and-statistics/398686/ltbcf-2016-final-508c-may-2018.pdf

Starsky has a similar philosophy toward validation – we begin with proper systems engineering design practices, adhering to industry standards where applicable, and then monitor leading indicators of performance and track outcomes, including harm. Starsky is focused on a risk-based approach where we aim to show that the risk of driving the automated system is comparable to or better than reasonable levels of risk posed by traffic already on the nation's highways. This assessment is iterative and improves as more data is created.

More specifically, we are using or creating minimal performance standards for our system behaviors. ISO 15622[16] is a performance standard and test procedure for ACC systems. One function of the Starsky system is ACC – the ability to hold a set speed while keeping a safe distance when traffic slows and return to the set speed when traffic speeds up again. ISO 15622 provides a set of tests and performance metrics for testing ACC on a vehicle.

Starsky can use this as a starting point when creating performance standards for this functionality.[17] Starsky creates test plans and go/no go criteria for each subsystem, which are then incorporated into a complete system plan for each milestone deployment. Go/no go criteria involve defining minimal performance requirements for each subsystem and then demonstrating that the system either does or does not meet those minimal requirements. This process removes pressure from individuals to green-light tests by implementing objective approval criteria to reduce risk. Having clear performance standards allows anyone with access to the test data to judge whether the system is meeting the required performance to allow it to be deployed.

The use of minimal performance standards is an important concept that helps focus our engineering roadmap. Because Starsky is creating a system aimed first at a very limited and specific ODD, our system design is easily captured by discrete system behaviors. Our goal is to narrowly define system behaviors in such a way that their performance can be measured quantitatively, and transitions between behaviors can be easily understood and evaluated. For example, "lane keeping" can be defined as a behavior with simple requirements:

- Stay in your lane;
- Hold a consistent speed; but
- Do not hit traffic ahead of you.

Because of Starsky's narrow ODD, each of these requirements can be quantified and tested. Starsky will only drive on highways with clear lane lines – so the Starsky system can be tested on these roads, and a performance metric can be created based on a ratio of undetected to detected lane lines, frame-to-frame consistency in identifying position relative to lane lines, cross-tracked errors as the

[16] ISO 15622-2018 Intelligent transport systems – adaptive cruise control systems – performance requirements and test procedures

[17] Starsky cannot adopt this directly as it was created for automobiles; trucks should behave more conservatively than automobiles because of the increased stopping distances for heavy trucks.

truck drives down the lane, etc. The Starsky system is not expected to drive without a teleoperator on more complicated roads, or on roads without lane lines, and the Starsky system can recognize when the lane marking quality is so degraded that accurate lane tracking is not possible. In a similar manner, Starsky can create metrics for the system's ability to maintain speed, detect surrounding traffic, and moderate speed to match traffic.

Starsky does not need to consider all the scenarios that developers planning to reach SAE J3016 Level 4 or 5 automation in a complex urban environment must be prepared to handle. For example, a truck on an interstate highway does not regularly encounter cross traffic, oncoming traffic, or traffic signals. The truck is unlikely to come across frequent bicycle, scooter, or pedestrian traffic. In the event that the Starsky system encounters these scenarios, they are automatically considered to be anomalous violations of our ODD, our safety architecture will be triggered, and the truck will respond accordingly. This difference in operating environment is a clear differentiator between an AV like the Starsky system, and AVs that are expected to be SAE Level 4 or 5 deployed in an urban driving environment.

These more complicated environments cannot be easily broken down into simple tasks that are easily analyzable. This is why many companies working to create these systems emphasize the importance of virtually simulating many, many scenarios as a validation tool. Some companies have simulated billions of miles and thousands of scenarios as a way to prove their systems are safe in these environments. Due to the difference in the quantifiability of our driving task and our use of a human-in-the-loop for certain decision-making processes, Starsky is not relying on virtual simulation to demonstrate the safety of the Starsky system. Instead, simulation serves as a supplementary tool for Starsky that provides a framework for testing individual modules, validating new features and improvements to the system, and conducting repeatable tests for system performance metrics.

Finally, Starsky tests the actual system on trucks that are driving on public roads with a safety driver in the cab. The safety driver is always a licensed, experienced professional commercial truck driver. The driver must be trained on the Starsky system before they are allowed to act as an actual safety driver in the cab of the truck. The Starsky system is driven to test actual performance under various conditions (with and without a trailer; with and without loads; heavy and light traffic; varying light and weather conditions) as a final qualification of new versions of system software. Ultimately, on-road driving is used as part of the process to whitelist specific routes to add to our ODD.

The goal of validation is to develop a statistically high-degree of certainty that the Starsky system is reliable enough to operate on a public road, monitored by a remote teleoperator, without a safety driver in the truck. Starsky uses a risk model based on techniques used for decades to justify launching unmanned systems and missiles at test sites.[18] Data from subsystem testing, actual

[18] The Range Commanders Council has been developing risk analysis standards for use at national test sites. See for example RCC 323-99 (Supplement), Appendix D. This method has been also used by the FAA to quantify estimates of the risk of flying unmanned aircraft and used in rulemaking proceedings for drone registration.

diagnostics, and on-road driving performance are used to make quantitative estimates of the probability of a failure, the probability that the Starsky system will detect that failure, and the probability that the system will then successfully execute a fallback behavior to achieve an appropriate MRC. The minimum threshold for Starsky to allow an unmanned test is to achieve "no unreasonable risk", which is equivalent to the level of risk posed by a large commercial vehicle operated by a human driver on public roads. In any instance where this confirmation threshold is not met, the Starsky system will not be deployed. In the early stages of deployment Starsky will add additional risk mitigations where needed to ensure that the system does not pose unreasonable risk to the public.

Starsky's validation strategy therefore relies on robust engineering practices that underlie the development of the system, as described in the section on System Safety. Unit tests are created to check the code against requirements with every build. The behavior of the Starsky system is then broken down into simple, analyzable segments, and a performance standard is created to make it possible to test each behavior as a subsystem of the larger system. Finally, the performance of the entire system is tested on-road, with a safety driver. These on-road miles are used to build a risk analysis to inform an objective decision regarding whether the system can be driven under a specific ODD with no safety driver aboard.

VII. Human Machine Interface

Starsky's strategy for the human-machine interface is shaped by our use-case. The Starsky system will only be operated by persons certified as professional drivers and trained specifically on the Starsky system. Humans will interact with the system in four ways: (1) as a driver in the truck (safety or manual); (2) as a teleoperator in a remote station; (3) as a member of the general public; or (4) as a public safety official.

Human drivers

The Starsky system is designed to allow a human driver in the cab to maintain the ability to manually operate and drive the truck at all times. The installed equipment does not interfere with or change how a human driver would operate the truck while sitting in the driver's seat.

The Starsky system includes an in-truck display that clearly indicates which operational mode is being executed at any given time. This display is connected to processors in the truck and will provide safety drivers or other individuals in the cab with information on system status, operational mode, current behavior, and built-in-test statuses. This will be the primary informational interface for any person in the truck, whether they are a safety driver or driving the truck manually.

The system includes a box installed on the dashboard to the right of the steering wheel with indicator lights and two push/pull switches. The indicator lights show the system status, including: engaged in autonomous mode, ready for engagement, unable to engage, or error detected in the

system. The switches allow a driver in the truck to engage autonomous mode, or to disengage the Starsky system completely with a single push of a large red button. This box will also issue audio alerts if an error is detected while the system is engaged.

The brake and accelerator pedals also include switches that allow the driver to disengage autonomous mode with a foot tap, much like common cruise control systems. Starsky has preserved this widely accepted method of disengaging vehicle automation.

The system also includes a 10-key keypad installed to the left of the steering wheel. This is used by a safety driver to mark events in the recorded data and set target speeds for the ACC. The driver can also push a key on this pad to disengage from autonomous mode.

Teleoperation

The Starsky teleoperation system is designed to allow a remote driver to safely operate and drive the truck from a remote location. As much as is practical, the teleoperation control interface mimics the actual truck control interface, complete with a steering wheel, brake and accelerator pedals, and other controls needed to perform the teleoperation task. The teleoperation system includes:

- Video screens that display the view from the driver's seat in the truck, including views of the mirrors and side windows, and views from cameras angled backwards from the sides of the truck.
- Telemetry displays that show the teleoperation-commanded positions of steering, braking, and acceleration, as well as the actual positions of each actuator.
- Telemetry displays of the current speed and target speed.
- A moving map indicating the position of the truck to help the teleoperator maintain situational awareness.

Public

The Starsky truck will operate in a public environment. Therefore, Starsky considers the interaction between the truck and the public (including public safety offers) to be part of the human-machine interface.

Like a manned truck, Starsky wants to ensure that the public has a method of contacting the company in case they are concerned by the Starsky system or a Starsky truck is involved in an accident. Starsky will clearly mark our trucks with a phone number connecting the motoring public with a responsible Starsky employee. The trucks will also have a unique ID number, so the public will be able to positively identify which truck system is being discussed.

Public Safety Officers (PSOs)

PSOs[19] will be able to use the same phone number to contact Starsky regarding any questions or concerns. For example, a highway patrol officer may see a Starsky truck pulled over on a shoulder. Using the unique ID and phone number, the officer will be able to easily contact appropriate staff at Starsky. We are working to develop a comprehensive law enforcement interaction guide that will contain best practices for officers engaging with a Starsky truck. This includes information regarding roadside and annual inspections, how to affirmatively disengage the automation system from affecting the truck in any way during an inspection, how the truck can be checked at weigh stations, towing trucks in the event of an accident, or otherwise pulling over a Starsky truck on the road. In advance of any milestone deployment, Starsky works with law enforcement communities regarding operational plans, proposed routes, and mitigation strategies in the event that a truck must be placed in a MRC.

VIII. Vehicle Cybersecurity

Starsky's cybersecurity program is a core component of our vehicle system security throughout the entirety of the design process. The cybersecurity program is designed to protect the ACMV from malicious attacks and promote public confidence in our systems. The scope of the cybersecurity program includes technical, administrative, and physical cybersecurity controls. Cybersecurity risk is different from other safety considerations in that cybersecurity envisions adversarial attacks rather than risks from design oversight or human failures, and is therefore inherently more difficult to detect or prevent.

Starsky's security umbrella includes threats posed by cyberattacks (exploits, "hacking"), social engineering attacks (phishing or obtaining passwords or other information that can be used in an attack), and physical attacks (breaking into a control room in an attempt to gain access to a truck). Starsky's risk assessment includes, but is not limited to, all of these possible attacks.

In general, Starsky is using the National Institute of Standards and Technology (NIST) Cybersecurity Framework to guide our efforts in this area. The NIST Cybersecurity Framework provides a thoughtful, deliberately defined set of topics and threats, and institutional actions categorized by robustness. Within the NIST Cybersecurity Framework, Starsky strives to follow the NIST Risk Management Framework. Through this continuous process we have designed a secure architecture and selected appropriate controls to mitigate risks inherent to the ACMV. Starsky monitors the threat environment and participates in industry cybersecurity groups to better understand, mitigate, and respond to new and emerging attack techniques.

Starsky works to apply appropriate cybersecurity best practices throughout the lifecycle of the design process and considers cybersecurity a critical factor in decision making. Our cybersecurity

[19] Public safety officers include anyone serving a public agency in an official safety related capacity, including law enforcement officers, firefighters, EMTs, and ambulance crew.

program is based on best practices from NHTSA, NIST, the Automotive Information Sharing & Analysis Center (Auto-ISAC), SAE, and others. Additionally, many system architecture design principles are borrowed from the aerospace and defense industries.

Starsky's cybersecurity system starts with an architecture that is designed to mitigate attacks and prioritize the safety of the system. Components of the system are designed to validate control commands and detect malicious inputs. Starsky evaluates our exposure to outside attacks by determining what aspects of the system could be accessed by outside parties. Access to the ACMV is granted through encrypted channels with robust authentication processes and is strictly limited to authorized users. Design reviews and code audits enforce a high standard of security through every stage of the development process. Extensive testing validates that the AMCV responds as intended in the event of an attack. Starsky scans for new vulnerabilities and performs simulated attacks against the ACMV to detect flaws. Finally, the Starsky system is outfitted with monitoring and detection systems, and Starsky has put response plans in place in the event of an attack. These measures are part of a continuing effort to improve cyber hygiene and implement applicable industry best practices as the cyber threat environment evolves.

In addition, Starsky's security strategy considers physical threats. We understand and prepare for potential physical attacks that could be carried out by malicious actors, including attempts to gain access to our trucks through forced entry to teleoperation centers or attacks against trucks traveling on the highway. The trucking industry has been the subject of increasingly sophisticated tactics and techniques used to hijack trucks or steal cargo. We consider these threats as part of our holistic approach to security and limit vulnerabilities by implementing controls and risk mitigation strategies accordingly.

IX. Crashworthiness

Starsky is installing our automation system as a retrofit kit to trucks that are commercially available on the market. We do not alter or disable any safety features and our system does not change the crashworthiness of the truck or change possible impacts of the truck on surrounding traffic.

X. Post-crash ADS Behavior

In the event of an accident, the system will choose an appropriate MRC depending on the state of the truck's sensors and systems, and traffic and road conditions. Where safe and possible, the system will choose to pull off the road onto an unoccupied shoulder and stop.

The Starsky system will have a remote teleoperator available at any time the truck is operating. After an accident, the teleoperator will be able to conduct an initial evaluation of the crash severity and take appropriate steps to alert authorities or public safety personnel. A teleoperator will be able to, for example, call 911 to report an accident. In addition, several states have specific statutory

requirements outlining the duties of ADS developers following an accident involving an AV. Starsky thoroughly evaluates post-accident procedures in each state where we operate and prepares to follow relevant requirements accordingly.

As long as the truck itself is not damaged badly enough to prevent driving and the teleoperation system is still functional, the teleoperator will be able to slowly drive the truck to a safe location where the truck can be parked. In all cases where the truck is drivable, a driver will be able to enter the truck, disengage the automated system, and manually drive the truck. The Starsky system does not change the truck's manual interface or prevent a human driver from driving the truck.

Starsky will evaluate the condition of the truck before any return-to-service after a crash. The truck itself will be inspected by a mechanic in the same way as any commercial vehicle before being returned to service.

The Starsky system is a collection of sensors and computers that are added to the truck. Each sensor can be evaluated and tested to make sure sensor quality has not been compromised by the crash. The equipment has built-in-test and self-monitoring capabilities that will be used to look for components that may have been damaged. Data is recorded on all systems in real-time, so Starsky can also review the recorded data to look for systems that may have malfunctioned, failed, or changed characteristics before, during, or after the crash. Starsky will perform a return-to-service process similar to the initial install, calibration, and check-out process that is used for every new system put into service.

XI. Data Recording

The Starsky system includes an onboard solid-state drive (SSD) array that records all activity of the system with sufficient fidelity to allow the data to be used to completely recreate all activity of the system. SSDs are a robust recording mechanism and the SSD should survive any crash where the entire truck cab is not destroyed or burned.

Data is recorded whenever the system is on and includes all camera and sensor data, vehicle on-board diagnostics (OBD) data, and all system diagnostics and warnings. The recorded data is used extensively for Starsky research and development. Starsky has created a rich infrastructure of data collection, storage, and analytics tools to fuel our ability to innovate.

Data is generally uploaded from each truck every evening to a Starsky cloud server. Starsky has tools running on the server overnight to look for and categorize problems so that our engineering team has a list of issues that can be reviewed the next day. These tools are subject to a continuous improvement process that helps the Starsky team efficiently review the huge amount of data that is generated each day. Starsky can use the recorded data to examine algorithm performance, recreate the

internal calculations of the system, and test changes to the algorithms. Each day of driving creates valuable data that can be used to benchmark performance and develop new automation methods.

In the event of an accident, this data recording would allow an engineer or investigator to recreate the events leading to the accident. This data includes video streams from seven cameras that provide views ahead, to each side, to each rear side, and from the driver's seat. The data also includes radar imaging to the front and sides. Additionally, the recorded data includes the system's determination of what objects are around the truck and what position and speed was measured for those objects, as well as truck system data such as speed, RPMs, control positions, and brake air pressure. The truck has redundant GPS systems and an inertial measurement unit (IMU) that can track the position and physical motion of the truck.

We believe that data sharing is an important component of government partnership. We have begun conversations with state officials regarding the potential value-add of Starsky's data and how to best create a platform to share data that may be useful to government stakeholders in a way that does not expose confidential business information.

XII. Consumer Education and Training

The Starsky system will only be operated by professional drivers with a commercial driver's license who have been specifically certified by Starsky. We view well-trained teleoperators as a critical component of enhancing highway safety. We believe our approach to automation offers the potential to transform the nature of the trucking workforce by offering well-paid, high-quality jobs in office environments that could reduce driver turnover. Our business model provides truckers an opportunity to utilize their experience in the industry while benefitting from enhanced flexibility, improved work life balance, and better working conditions. Starsky's vision for the future includes a robust labor force development, training, and certification pipeline that builds the technical skills required for teleoperators and prepares the American labor market for the changing nature of work in the 21st century.

In addition, consumer and public education will be critically important as AVs are introduced to U.S. roadways. Starsky has concerns regarding the reaction of the motoring public when they find themselves on the highway next to an ACMV without a human operator in the driver's seat. Today, motorists often engage in aggressive driving behavior around legacy CMVs, including excessive speeding, tailgating, and erratic lane changes without signals, all while often misjudging truck speed. This behavior could be potentially exacerbated in relation to ACMVs, or cause driver distractions when motorists see trucks without a human in the cab. Starsky is considering potential solutions to increase awareness surrounding ACMVs and educate the public on how to properly interact with trucks without a human driver. In the end, a strong element of public outreach will be required to build public acceptance as ADS reach some critical mass of traffic on the road.

Starsky views our press operations and public relations efforts to be part of this outreach effort. Positive stories about the future of the trucking industry, emphasis on improving working conditions for drivers, interest in the high-tech aspects of AVs, and details regarding our approach to safety – these are all messages that can help the public understand the value and benefits of Starsky's work. Due to Starsky's focus on route-specific ODDs, we can take additional steps, including working with government officials and journalists specific to the areas where Starsky will operate. Journalists can be helpful in educating the public by publishing stories about driverless trucks that may be seen in an area, helping the public understand that these deployments are being executed safely and responsibly, increasing public interest in ACMVs, and generating local pride in their region's role in advancing life-saving next-generation transportation technologies. Our extensive outreach efforts with local governments prior to deployments help keep relevant officials informed, so when constituents call to question the presence of an unmanned truck, government stakeholders are well-informed and not caught unaware. Local officials are extremely helpful partners in explaining ACMV safety to their constituents and increasing awareness regarding the importance of facilitating ACMV development to improve highway safety.

In addition, we are working on the development of more comprehensive public education campaigns that seek to build consumer trust in the Starsky system, teach communities how to interact with our trucks, and prepare the public for ACMV deployment.

XIII. Federal, State, and Local Laws

Starsky works to ensure that we are compliant with applicable federal, state, and local laws wherever we operate. Today, the rapid development of ADS technology is outpacing the development of corresponding regulation, and in many cases, ACMVs have been overlooked or deliberately excluded from state and federal regulations. Therefore, Starsky has employed a proactive approach to government affairs and self-regulates to go above and beyond what is required by law. Starsky works closely with relevant legislative, regulatory, and law enforcement entities before we deploy in a given jurisdiction.

In some states, for example, AV statutes do not outline formal certification, approval, or reporting processes for testing or deployment, and simply require ADS developers to meet certain operational, insurance, and registration criteria. Despite a lack of formal approval processes, Starsky actively engages relevant authorities and stakeholders, including law enforcement communities, to ensure that affected state and local officials are well aware of our operating plans and system capabilities in advance of deployment. In most states, government agencies have designated personnel who serve as points of contact with jurisdiction over ADS-related policy and enforcement.

In addition, Starsky regularly consults law enforcement agencies concerning the development of best practices to allow interactions between public safety officials and Starsky trucks. This includes mechanisms for day-to-day law enforcement interaction, such as a unique telephone number and ID

displayed on our trucks to allow officers to connect with an appropriate contact who can respond to requests and facilitate inspections.

Starsky also exercises strict compliance with all requirements mandated by Federal Motor Carrier Safety Regulations (FMCSRs), including hours-of-service, inspection requirements, and medical qualifications for drivers. As FMCSA considers changes regarding the applicability of these requirements to ACMVs, we will continue to work with the agency in future rulemaking proceedings to strike the right balance between preserving existing operational requirements for truck drivers, which we believe are critical to safe CMV operations, and changing regulations to support ACMV deployment. In the interim, Starsky will continue to engage the FMCSA to seek regulatory relief under existing exemption authority if necessary, understanding that a request for a waiver or exemption must demonstrate Starsky's ability to meet equivalent safety standards.

Starsky will not deploy our systems unless we have explicitly confirmed our compliance with applicable laws and norms in a given jurisdiction. While not required by states in which we operate, Starsky has taken additional steps to secure outside counsel to develop written legal analyses outlining our compliance with relevant statute and regulation before we deploy in a given state. During our route feasibility whitelisting process described above, Starsky thoroughly vets every route. Through this process, we confirm that our operations will comply with federal, state, and local requirements, such as traffic laws, speed limits, noise ordinances, and move-over laws. Starsky monitors changes in regulation in jurisdictions where we operate and will take steps to comply accordingly.

To date, Starsky has found federal, state, and local partners to be immensely helpful resources regarding our compliance with applicable regulations. Starsky sees proactive government outreach and fostering close relationships with regulators and elected officials as critical to the deployment of ADS technologies. Through these relationships, Starsky works to ensure our industry is allowed to continue to innovate while the safety of the motoring public is monitored by the appropriate authorities.

XIV. Appendix

Notes on Risk and Safety

The terms "risk" and "safety" can be defined and quantified. Starsky uses the following definitions.

Safety is the absence of *unreasonable risk*.[20]

[20] This definition for safety is derived from the definition in ISO 26262.

Risk is the combination of a hazard (harm) and a likelihood of exposure. For automobiles and trucks, this is commonly characterized in terms of "fatalities per 100 million miles", or "fatalities per fleet hour" for airplanes.[21]

There is no such thing as "perfect safety" or a vehicle that will never cause harm. We are surrounded by technologies that cause harm every day, including automobiles and trucks. NHTSA data shows that more than 37,000 people were killed in automobile accidents in 2017[22], or approximately 100 deaths per day.

Unreasonable risk is defined by society. This is the defining line between risks that are generally accepted, and risks that are not deemed acceptable by society at large as part of the adoption of a practice or technology. This acceptance threshold generally changes over time as technology evolves and people expect safer products. A risk goal in general aviation aircraft, for example, is one fatality per 100,000 flight hours. This statistic is tracked using actual fleet data[23], and also used as a goal when the certification of new equipment or aircraft is being considered.[24] For general aviation, this defines the threshold between reasonable and unreasonable risk. Automobiles have demonstrated a remarkable evolution in risk with a corresponding change in the public's expectation of unreasonable risk.[25] Over the better part of the last century, automobile accidents have fallen from a fatality rate of 24 deaths per 100 million miles to one death per 100 million miles. This has been achieved through a combination of changes in culture (views about driving under the influence of alcohol or drugs), changes in laws (maximum speed), infrastructure (improved highways and urban traffic controls), and equipment (seat belt, airbags, and anti-lock brakes). Federal regulation has kept pace with the expectation of reduced risk: today, it is impossible to legally sell a new automobile in the U.S. marketplace without seatbelts and airbags.[26]

It is common for new technologies to be benchmarked against existing technologies. It is assumed that the risk posed by the existing, deployed technology is already accepted by society, and therefore the risk posed by that technology is not unreasonable. We should expect that society justifiably has a stake in the deployment of new technologies. New technologies should provide some benefit, and this benefit provides the push to deploy the new technology. Setting the risk threshold too conservatively may deprive society of this benefit unnecessarily. Setting the risk threshold too readily may expose society to additional harms not yet justified by the benefit. Often, both the potential

[21] This definition of risk is used across multiple industries and is the basis for risk analyses used in regulating medical devices, aircraft, and automobiles in one form or another.
[22] https://crashstats.nhtsa.gov/Api/Public/ViewPublication/812603
[23] The National Transportation Safety Board collates statistics on aviation accidents and fatalities. https://www.ntsb.gov/investigations/data/Pages/AviationDataStats2015.aspx
[24] AC 23.1309 "System Safety Analysis and Assessment for Part 23 Airplanes", figure 2, contains explicit design goals for risk for various types of aircraft, combining severity of hazard and likelihood of occurrence. https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_23_1309-1E.pdf
[25] Wikipedia summarizes statistics from a variety of sources on motor vehicle fatality rates. https://en.wikipedia.org/wiki/Motor_vehicle_fatality_rate_in_U.S._by_year
[26] "Frontal airbags have been standard equipment in all passenger cars since model year 1998 and in all SUVs, pickups and vans since model year 1999." https://www.nhtsa.gov/equipment/air-bags

benefits and harms are difficult to assess with any accuracy greater than an order of magnitude. What is considered an "unreasonable risk" is not something that can realistically remain static for all time. The level will change with time, as society gradually raises the expectation of what is reasonable and unreasonable.

Starsky will benchmark our ACMVs against the safety record of manned trucks. FMCSA has conducted studies regarding crashes involving trucks and has identified key driver-related factors that can contribute to accidents. Human drivers can suffer from fatigue or boredom which can lead to bad judgements and inattention on the road. For example, a truck driver that has been delayed on a long drive might decide to speed when traffic eases to make it to a stopping point before reaching their hours-of-service limits; or after spending many hours on a monotonous, straight interstate, their attention may wander. These human weaknesses are exactly the areas where the Starsky system is designed to be strong. A long, straight interstate is easy for the automated truck to drive, and the truck does not get tired or worry about where to sleep each night. Starsky will take advantage of the expertise of human drivers to handle the complicated traffic situations off-highway, where automated systems have difficulty driving safely. The automation will help precisely where humans are most susceptible to unsafe driving, and professional drivers will help precisely where automation is most susceptible to unsafe driving. Together, the combination of automation and professional drivers has the potential to make the Starsky system significantly safer than legacy heavy-duty trucks on U.S. highways.