




DATA ANNOTATION AND AUTONOMOUS VEHICLE

SAFETY



HOW SAFE IS SAFE ENOUGH?

How safe is safe enough? When the comparison is between a car driven by a computer versus a person, the practical answer depends on competing perspectives.

Purely from a statistical standpoint, Autonomous Vehicle developers have a clear target. The United States alone sets a grim annual benchmark. According to the National Highway Traffic Safety Administration, there are more than six million car crashes every year. More than 36,000 Americans die in those crashes, with another 2.5 million ending up in hospital emergency rooms.

Considered globally, the figures are even more sobering.

What if a switch to fully Autonomous Vehicles cut those figures in half, could AV developers declare victory? Perhaps from the standpoint of cold logic, but those involved in various AV initiatives would concede that the more telling dynamic is consumer confidence – as evidenced by their resulting willingness to buy a fully Autonomous Vehicle, or even step inside one. Developers grimace at the thought of having to defend the idea of tens of thousands of deaths from collisions involving driverless cars.

A 2018 Rand Corporation report, “Measuring Automated Vehicle Safety,” alludes to the conflict between the need for empirical data for AV developers seeking to advance the state of the art, and the determination of consumer safety regulatory bodies to resist what they generally regard as avoidable risk.

“In the United States – and elsewhere, to some degree – the emergence of AVs has been associated at least implicitly with the view that some exposure to risk and uncertainty about this risk must be accepted in the short and medium terms to see the long-term benefit of AVs,” the report notes. “(But) consultations for this project showed how that view, held by AV developers and by associated industrial and research partners, is not held by the safety advocacy community, which champions clearer communications about risk and more conservative efforts to at least minimize risk and preferably eliminate it.”

THERE'S BEEN A FASCINATION WITH AUTONOMOUS VEHICLE TECHNOLOGY FOR DECADES... AS DEPICTED IN THIS AD FROM 1950 THAT ENVISIONED THE CAR OF THE FUTURE.



ELECTRICITY MAY BE THE DRIVER. One day your car may speed along an electric super-highway, its speed and steering automatically controlled by electronic devices embedded in the road. Travel will be more enjoyable. Highways will be made safe—by electricity! No traffic jams... no collisions... no driver fatigue.

POWER COMPANIES BUILD FOR YOUR NEW ELECTRIC LIVING

Your air conditioner, television and other appliances are just the beginning of a new electric age.

Your food will cook in seconds instead of hours. Electricity will close your windows at the first drop of rain. Lamps will cut on and off automatically to fit the lighting needs in your rooms. Television "screens" will hang on the walls. An electric heat pump will use outside air to cool your house in summer, heat it in winter.

You will need and have much more electricity than you have today. Right now America's more than 400 independent electric light and power

companies are planning and building to have twice as much electricity for you by 1965. These companies can have this power ready when you need it because they don't have to wait for an act of Congress—or for a cent of tax money—to build the plants.

The same experience, imagination and enterprise that electrified the nation in a single lifetime are at work shaping your electric future. That's why in the years to come, as in the past, you will benefit most when you are served by independent companies like the ones bringing you this message—America's Electric Light and Power Companies*.

*Names on request from this magazine.

Yet as the report goes on to observe, "Given that clear measurement of safety is not likely to be obtained without exposing the public to AV risk on public roads, incremental increases in public AV use are likely to increase our knowledge about safety. The incremental approach can limit risk without stifling the technology. (Yet) we recognize that the long-term, big-picture view based on confidence in the merits of AVs is shaken every time there is a crash."

That's exactly what happened the year the report was released.

A single fatal encounter in 2018 involving an Autonomous Vehicle owned by the ride sharing company Uber (See: Uber, sidebar page 11) forced the company in the immediate aftermath to abandon open road tests and ultimately set back the company's AV development program by years. Even the prospect of a few dozen accidents a year causes the AV development community to cringe – regardless of whether or not that would represent a significant improvement in road safety.

The Uber incident shook the Autonomous Vehicle development community by exploding assumptions on algorithm capabilities in real world environments. Many AV developers, for example, eliminated or severely restricted the use of public roads for system testing and instead moved to internal test tracks. But more importantly, they reassessed the level and complexity of algorithm training required to substitute for a human driver. And they began investing heavily in more diverse and detailed sets of training data.

The timing and intent on demonstrating the industry's ability to focus on safety concerns has proved productive. A little over two years after the Uber incident, the United States Department of Transportation's National Highway Traffic Safety Administration (NHTSA) has issued a notice of proposed rulemaking addressing safety policy for Autonomous Vehicles, as well as a pilot web-based project to coordinate public visibility into Autonomous Vehicle testing.

The NHTSA web portal, announced in June 2020, includes data from nine AV system developers and car manufacturers – Beep, Cruise, Fiat Chrysler Automobiles, Local Motors, Navya, Nuro, Toyota, Uber, and Waymo – tracking testing activity in eight states. The web portal includes mapping tools showing testing locations at the local, state, and national levels, as well as testing dates, frequency, vehicle counts, and routes.

The initiative also reflects a tacit agreement between the industry and government regulators that Autonomous Vehicle development can't be solely limited to closed testing tracks. There is no substitute for real world testing, as a 2018 Automobile Association of America study found.

AAA tested the driver assist systems in cars from five manufacturers (Ford, BMW, Kia, Subaru, and Cadillac), and while the Advanced Driver Assist systems generally "performed according to expectations derived from the owner's manuals during closed course evaluations," tests on public roadways, including highways 101 and I-5 in California, resulted in one ADA "event" roughly every eight miles. AAA defined

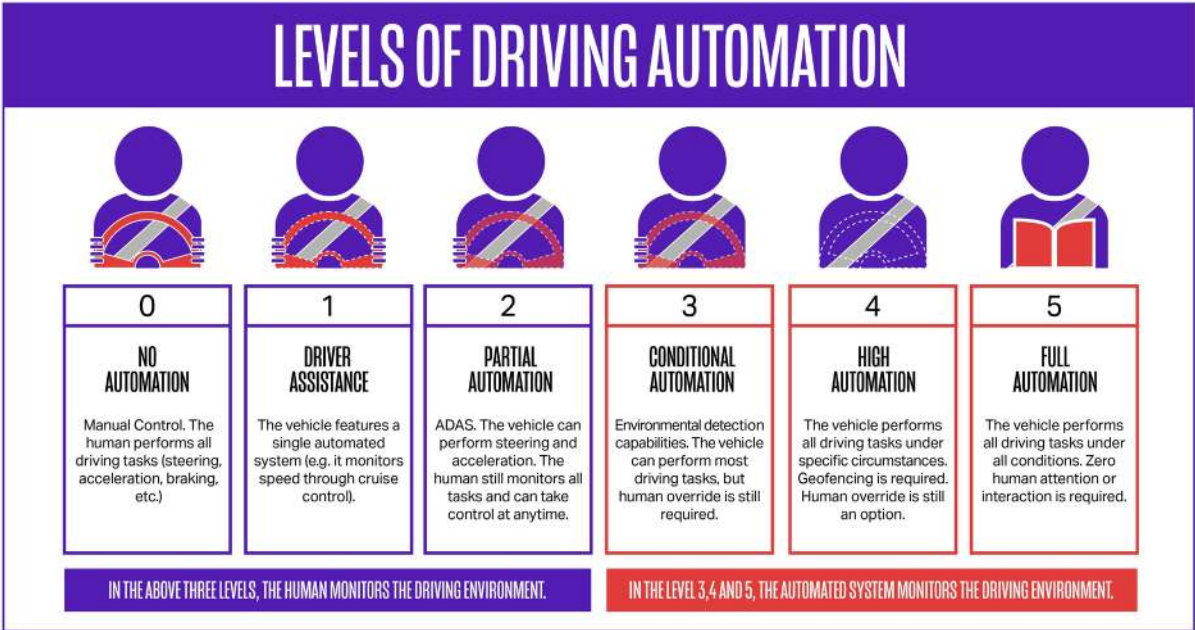
an event as a lane departure, erratic lane positioning, or a failure to slow or stop to avoid a near collision. Overall, AAA gave the systems reasonably positive – but not quite passing – grades.

“Nothing contained in this report is intended to directly imply (that) evaluated ADA systems are inherently unsafe,” the report states. “In totality, this work is intended to illustrate potential system inadequacies in the context of scenarios realistically encountered in a naturalistic environment. Additionally, all test drivers expressed some degree of distrust in each evaluated system and questioned the overall usefulness of integrating adaptive cruise control and lane-keeping functionalities.” Translation: All of the systems needed more training data to fill in gaps in the algorithmic knowledge. And to a great degree, system developers went back to work.

The increased attention to Autonomous Vehicle safety issues over the past 2-3 years also reflects the evolution in car and system development – and a bit of irony in public perception. Car manufacturers were quick to (rightly, for the most part) position early industry successes in AV Level 1 and Level 2 so-called “Driver Assistance” systems as victories for enhanced driver safety. Systems designed to keeping a moving car within its lane, and to detect and prevent imminent crashes through emergency braking systems were designed specifically to improve safety.

But public perception shifted as car manufacturers moved on to Level 3 and Level 4 system development, where the aim shifts from assisting a driver to fully automated operations, with the driver there to take over only in an emergency. With a computer now driving the car, many consumers do not accept any margin for error – certainly not at highway speeds.

As developers have conceded in conversations on the subject, consumers are not about to demand that manufacturers pull a car model off the road if a driver assist system fails to correct for human driver error. But if an accident occurs with a computer at the wheel, such as an Autonomous Vehicle failing to recognize a “Do Not Enter” warning sign? It will be the machine on trial, with a swift, punitive judgement.




In theory, a Level 5 fully Autonomous Vehicle ought to be safer than one driven by a person. An AV sensor array of radar, LiDAR and cameras won't be distracted by a roadside billboard, or ringing cellphone. AV systems don't suffer from driver fatigue, or uncertain reflexes.

But a human driver brings a rich set of experience and direct knowledge about the surrounding environment. And as the many new AV safety standards suggest, AV system developers still face a huge challenge contemplating all of the scenarios an Autonomous Vehicle must be prepared for – and another hurdle developing, processing and analyzing all of the training data required to convert those scenarios into computing knowledge.

The irony is no one needs to teach even a new driver that wet leaves represent a potential hazard – but where some drivers might ignore the risk a properly trained AV algorithm never will.

AV developers could, of course, limit risks by reducing the performance parameters for Autonomous Vehicles. Instead of programming a Level 5 Autonomous Vehicle for posted highway speed limits, developers could create a "cautious driver" profile that limits highway performance to 40 MPH. But the entire industry realizes that blunting utility as an artificial means of managing risk is a double-edged sword of sorts; having Autonomous Vehicles create five-mile highway traffic jams out of extreme caution is in the end a self-defeating proposition.

The good news is there are algorithmic methods of managing risk while still maximizing vehicle performance. There are already mathematic formulas in place to set a safe distance between vehicles, for example, one where specific vehicle braking performance (a Porsche will always out-brake an economy sedan) can be plugged in to account for mechanical as well as computational capabilities.



Safe Distance Formula

$$d_{\min} = L + T_f [v_r - v_f + \rho(a_a + a_b)] - \frac{\rho^2 a_b}{2} + \frac{(T_r - T_f)(v_r + \rho a_a - (T_f - \rho)a_b)}{2}$$

- L is the average length of the vehicles
- ρ is the response time of the rear vehicle
- v_r, v_f are the velocities of the rear/front vehicles
- a_a, a_b are the maximal acceleration/braking of the vehicles
- T_f is the time for the front car to reach a full stop if it would apply maximal braking
- T_r is the time for the rear car to reach a full stop if it would apply maximal acceleration during the response time, and from there on maximal braking

The above formula calculates the safe longitudinal distance between the rear vehicle and the front vehicle.

EVOLVING PRIORITIES FOR AUTONOMOUS VEHICLE ANNOTATION

As Autonomous Vehicle system development has advanced from one level to the next, from partial to higher degrees of automation, it was inevitable that the underlying focus on the part of developers would expand from improving the reliability of basic functional capabilities to operating safely within complex on-road scenarios. That has been reflected in subtle but clear changes in data analysis and annotation priorities.

Where the early development focus had been almost solely on expanding and deepening algorithmic capabilities to accurately incorporate “ground truth” data – the ability of an Autonomous Vehicle to recognize and differentiate between static inanimate objects on one hand, such as mail boxes, street lights and the like, and dynamic objects (everything from other cars to human beings, and animals) – there is at least as much of a priority now on more complex real-world scenarios.

Most are triggered by the failsafe option in AV testing: takeover events, where the AV algorithm is unable to cope with an event or series of events and hands over operation back to a human test driver. Before Autonomous Vehicle developers can deliver Level 5 fully Autonomous Vehicle operations, they have to address and solve those “brain” freeze moments (See: Uber, page 10) when the algorithm becomes so confused or indecisive that it is unable to safely maneuver the vehicle.

Most evolved AV algorithms are able to recognize, for example, a bicycle and a human on or alongside the road. But the shape and nature of each of those objects are fundamentally altered when they merge into a human riding a bicycle. Until the algorithm has processed sufficient annotated visual data to recognize those objects both individually and together (and the latter could include a person walking a bike instead of riding it), it will lack the intelligence to “see,” process, and react to those common situations.

Heightened intelligence comes both from the quantity and quality of the training data fed through an AV algorithm. In practical terms, developers have begun to rely on far more detailed object taxonomies in order to begin to mimic the human understanding an experience. Where in the past some AV system developers might have been satisfied with a catchall designation of “vegetation” to categorize anything that grows, there’s a greater priority on segmentation – to trees, bushes, grass, flowers, for example – where each object presents different characteristics (and different potential threats).

Basic ground truth data annotation is a fairly straightforward affair – the application of bounding boxes, polygon analysis, and the application of semantic segmentation to teach the algorithm what another car looks like, or a mailbox, or a tree. What sets a good skilled and experienced data analyst apart from a beginner (and sets a top tier data annotation vendor apart from more pedestrian firms) is the ability to perform the type of problem-solving data analysis that addresses all imaginable Autonomous Vehicle safety issues.

Typically, the data analyst will examine each takeover event from the perspective of two scenarios, one of which is what actually happened after the test driver took over control of the vehicle. The second involves a so-called “ghost car,” the computer simulation vehicle as it would have continued under the control of the algorithm. It’s the job of the data analyst working with the algorithm developers to determine (1) why the human driver had to take over, and (2) what would have happened had the onboard computer remained in control of the vehicle. Would there have been an accident – and, if so, how severe would it have been?

From the standpoint of division of labor, an iMerit data analysis team would, for example, work to identify the object (or combination of objects) that caused the car to hand off operations to the driver. Developers then take the information and use it to further refine the algorithm, typically by sending more training data for annotation.

Sometimes the event post-mortem investigation identifies an overreaction rather than a near-collision. If the algorithm is not familiar with leaves, for example, (perhaps all of the training data was gathered in the summer) it might “panic” when the vehicle’s sensors pick up leaves on the road. The algorithm might not understand if the leaves represent a dangerous obstacle or a benign presence. Are there enough of them – and are they wet enough – to present a hazard to the vehicle. Or are they outdoor decoration the algorithm can note and then ignore?

The complexity of the development task ratchets up further when the algorithm is tasked with predicting the movement of other objects, whether living or inanimate. First the algorithm must ingest the data captured by the Autonomous Vehicle’s array of sensors and compare that information with what it has learned from training data defining different objects and their behavior patterns (if any). It then needs to recall behavior patterns based on the nature of each object and then predict how it might move – and (crucially) how those movements might intersect with the vehicle.



Level 5 Autonomous Vehicles must detect, analyze, and react to unpredictable environments. Source: AAA

THE NEED FOR AUTONOMOUS VEHICLE STANDARDS

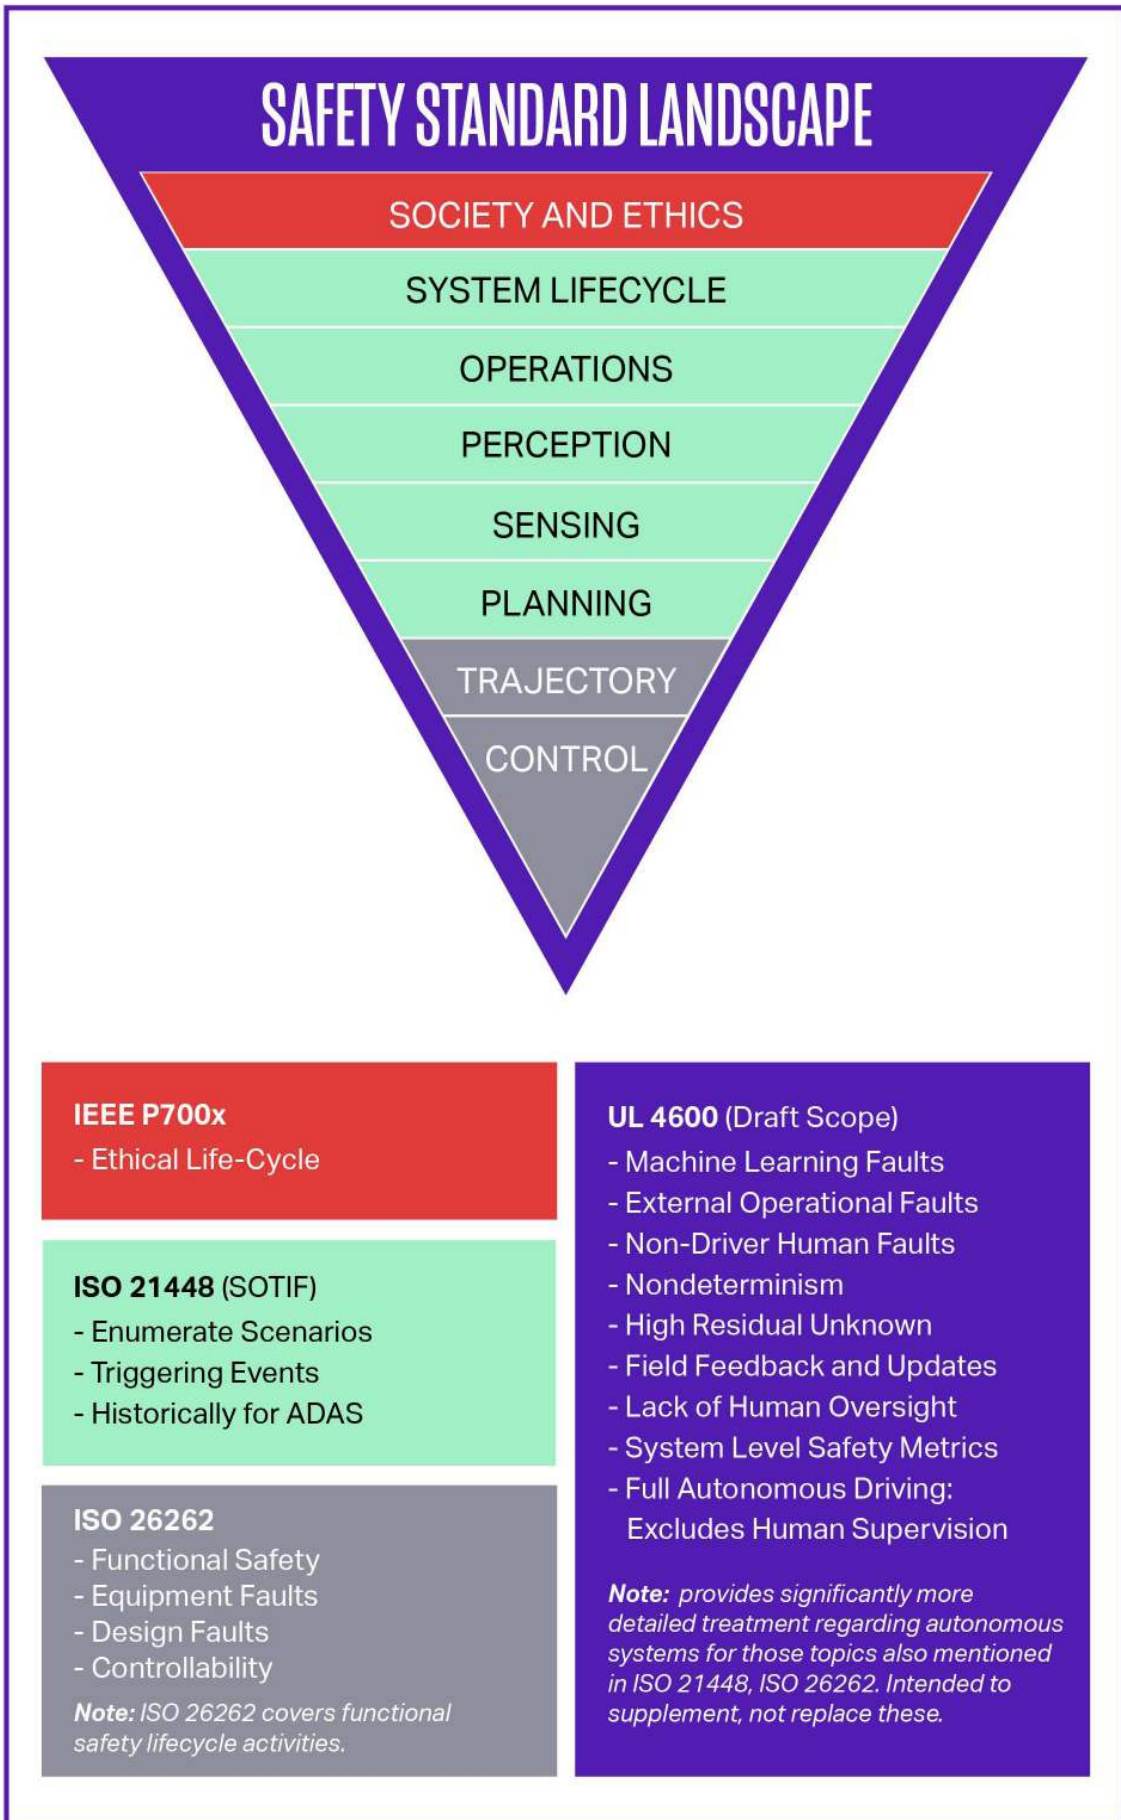
Prior to 2020, Autonomous Vehicle development, focused on improving the safety of AVs, took place in a regulatory and standards vacuum, with most AV developers taking an independent approach to safety metrics and system requirements. That's not really surprising, given the cutthroat, winner-takes-all nature of technology development. But the usual suspects of complex systems issues, from the acknowledged need to ensure interoperability between different systems and subsystems, and the inevitability of multiple AV system vendors that are supplying the leading car manufacturers, has pushed the industry towards a greater degree of coordination and collaboration than automobile manufacturers had previously wanted to embrace.

By the time the first Level 5 fully Autonomous Vehicles go on sale (and few will hazard a guess on that date) there could be as many as a dozen different industry and governmental safety standards and certification programs. Among the standards getting the most attention are the International Organization for Standardization's (ISO) 21448 and 26262 as well as Underwriter Laboratories' (UL) 4600 (see UL sidebar, pages 12-13).

The flurry of activity around AV safety and standards development was as inevitable as it was necessary. Government regulators, standards setting bodies and the automotive industry, itself, all understand that the transition from human drivers to driverless cars will necessitate a reassessment of safety issues, ranging from how manufacturers and system developers architect their solutions to all of the operational scenarios that would have to be judged through a completely new contextual lens.

ISO 26262, an offshoot of an earlier ISO standard for automotive systems, was developed to address the functional safety of automotive electrical systems, including the onboard computers, sensors and other systems in Autonomous Vehicles. As a byproduct of that focus, the standard will likely be used to guide the identification and selection of algorithm training data. That is certainly the expectation for ISO 21448, better known as SOTIF (Safety of the Intended Functionality), which specifically addresses Autonomous Vehicle development. The standard in part contemplates the need for scenario-based algorithm training data to account for all of the significant functional situations an Autonomous Vehicle onboard computer might encounter.

Unlike some of the other standards that apply to automotive safety, SOTIF does not address system failures. Instead, it was written to encompass so-called "safety violations" that occur without a system failure – such as when the development of an AV algorithm fails to contemplate all of the operational scenarios an Autonomous Vehicle might encounter in a real-world environment. As such, the standard is expected to be a guide of sorts – and an essential companion to UL 4600 – in formulating all of the algorithm training scenarios developers will need to address through appropriate training data.



SAFETY STANDARD LANDSCAPE

SOCIETY AND ETHICS

SYSTEM LIFECYCLE

OPERATIONS

PERCEPTION

SENSING

PLANNING

TRAJECTORY

CONTROL

IEEE P700x

- Ethical Life-Cycle

ISO 21448 (SOTIF)

- Enumerate Scenarios
- Triggering Events
- Historically for ADAS

ISO 26262

- Functional Safety
- Equipment Faults
- Design Faults
- Controllability

Note: ISO 26262 covers functional safety lifecycle activities.

UL 4600 (Draft Scope)

- Machine Learning Faults
- External Operational Faults
- Non-Driver Human Faults
- Nondeterminism
- High Residual Unknown
- Field Feedback and Updates
- Lack of Human Oversight
- System Level Safety Metrics
- Full Autonomous Driving:
Excludes Human Supervision

Note: provides significantly more detailed treatment regarding autonomous systems for those topics also mentioned in ISO 21448, ISO 26262. Intended to supplement, not replace these.

THE UBER CRASH: IGNORANCE CAN BE DEADLY

At 9:58 on the evening of March 18, 2018, 49-year old Elaine Hertzberg was walking her bicycle across a four-lane road in Tempe, AZ, as an autonomous test vehicle operated by the ride share company Uber approached. The events of the next few seconds would result in human tragedy, and a re-examination of what it takes to put a fully Autonomous Vehicle on the road.

It took 19 months for authorities to piece together a clear picture of the accident and its causes. In the end, it became clear that what killed a pedestrian that night, and set back Uber's ambitious AV test program, was an all-too common combination of human and machine failure – ultimately rooted in a gap in AV algorithm development. In that event, the collision provided a painful lesson to the industry: the benchmark for AV safety sits well above the accepted norms for human-controlled cars. Where public sentiment is concerned, there is no room for failure.

The initial human failure – no fault of the Volvo 2017 XC90 SUV outfitted with multiple sensors (forward and side-facing cameras, radar, LiDAR, navigation sensors, and an onboard AV computer) – lay in the hands of the “failsafe” test driver placed in the driver's seat (to take over in the event of a safety system failure). The driver, according to the accident report, was streaming Hulu up until the moment the car collided with Hertzberg.

Ironically, the Volvo's radar and LiDAR sensors detected Hertzberg in plenty of time to either brake or swerve – six seconds before impact, at a speed of 43 MPH. But the training data for the AV algorithm had not accounted for the possibility of a pedestrian outside a crosswalk, and the system cycled through potential object identification conclusions (unknown object, vehicle and, finally, a bicycle) as the vital seconds ticked down.

Hertzberg, according to cameras in the Volvo, was wearing dark clothes when she crossed in an unlighted section of the road, and there were no reflectors on her bicycle. While the Volvo's radar and LiDAR sensors would still have detected her in time to initiate emergency braking, those conditions, coupled with gaps in the AV algorithm's training data, might have contributed to the system's deadly confusion.

Just 1.3 seconds before impact, the car's onboard computer attempted to slam the brakes – but to no avail. Uber had modified the car to deactivate the factory emergency braking system when the car was being controlled by the aftermarket onboard Autonomous Vehicle computer – ironically, “to reduce the potential for erratic vehicle behavior.” Finally, less than a second before impact the safety driver took the steering wheel. By the time she hit the brakes, a second later, it was too late.

In the months following the accident, Uber would add training data to account for jaywalkers, the company later told NTSB investigators. But news reports following the accident cited the company's own internal documents suggesting that – with vehicle test driver takeover events occurring roughly every 0.8 miles – the AV onboard computer was not nearly ready for primetime, nor driving on public roads.

PRODUCT SAFETY ICON ADDRESSES AUTONOMOUS VEHICLES

Of all the emerging standards helping to guide Autonomous Vehicle car development, the standard that is garnering the most attention is a recent release from Underwriters Laboratories. Although based in the United States, UL defines and certifies product safety for a global audience. Its foray into AV development brings a gravitas that promises to command the attention of car manufacturers and other AV system developers around the world.

From the inception, UL expected that its UL 4600 standard for Safety for the Evaluation of autonomous products would work in conjunction with other AV-focused safety standards, including from the FAA, IEEE, and ISO. But unlike some of those other, more tech-centric AV-related standards, UL 4600 addresses “safety principles” for evaluating fully autonomous products rather than specific technologies and methods. UL 4600 doesn’t tell a car manufacturer how to build a safe driverless vehicle; instead it instructs them on how to make the case for the safety of their cars. Cars are the first intended use case, but the standard applies to other autonomous products, including everything from mining equipment to drones.

“The safety case,” notes the just released standard document, “includes a structured set of claims, arguments, and other evidence supporting the proposition that an item (a vehicle plus all the other support contributing to safety) is acceptably safe for deployment.” The standard, it notes, “does not define a process, but rather puts forth assessment criteria to determine the acceptability of a safety case.”

Those safety cases are built largely around the many scenarios an Autonomous Vehicle will have to navigate – literally as well as figuratively. For example, “(a) pedestrian appearing suddenly in (the) roadway after crossing in front of a public bus discharging passengers.” Or “(a) human ability to command traffic law violations such as going through a red light due to perceived danger to occupants if remaining stopped at (a) light.” Or even “(an)other vehicle skidding into (an) intersection due to ice, collision from the rear due to close following distance, (a) vehicle being pushed into an intersection due to collision from the rear, distracted driving behaviors, (or) impaired driving behaviors.”

The standard is expansive enough to consider even events not typically envisioned by car manufacturers, such as aggressive behavior by pedestrians that go beyond stepping off a curb without first looking both ways. This category includes unusual yet conceivable events a human driver would have to assess and react to, including the category of “vehicle harassment by non-occupants,” such as “defiant jaywalking, squeegee punks... (or) intentional injury seeking behavior” such as a “pedestrian intentionally (jumping) in front of a vehicle.” It even contemplates scenarios where a human driver would have to quickly assess and take action based on overtly threatening behavior rather than merely displaying quick reactions. In that category are “situations such as malicious pedestrians exploiting designed (AV) safe item behaviors to nefarious purposes, such as blocking the movement of a vehicle in order to entrap the occupants to attack them.”

In those abhorrent behavior scenarios, UL 4600 does not call for what would amount to a superhuman reaction on the part of a car's onboard computer. Instead, the standard notes, "this is not a requirement to absolutely prevent impact with a pedestrian who is willfully attempting to be injured. Rather, this identifies a scenario that should be considered in hazard analysis, and in particular might reduce the credit that can be taken for humans avoiding risky situations."

While UL 4600 does not mandate solutions to each of those scenarios, or even minimum milestones for AV system development, such as specific data annotation and analysis, it does back into that territory in the standard's Section 8.5 Machine Learning and AI Techniques. For a car manufacturer to win UL 4600 certification, it will have to make the case that its onboard AV system contemplates each of those above scenarios (and more) and that its algorithm is fully capable of detecting, assessing, and then reacting to those sets of conditions and events.

"The safety case (put forth by a car manufacturer) shall argue that any Machine Learning based approach and other AI approaches provide acceptable capabilities," UL 4600 notes. In that respect, it calls for explanations demonstrating that "the type and quantity of data used for Machine Learning training and testing ensure acceptable performance across the entire (Operational Design Domain), addressing at least the following points: (1) (that the) safety related aspects of the ODD are substantially represented by the data, and (2) evaluation metrics account for required risk mitigation." The ODD for an Autonomous Vehicle can be defined in terms of geography, weather, lighting, roadway markings, previous experience on that particular roadway, and other environmental and experiential factors.



In plainer language, the standard offers a working scope for AV Machine Learning training data: "(That) events with a low probability but high severity for failure are properly handled, despite being a very small fraction of expected real world data." In doing so, UL might have just caused AV system developers to reassess the quantity, quality, and diversity of the training data they will require to pass perhaps the most fundamental test for consumer acceptance: a car that not only reacts as well (if not better) than a human driver to everyday road hazards – but even to the dangers contemplated by the fevered imagination of a Hollywood thriller.

ANSYS	Beijing Research Institute of Automation for Machinery Industry	Intel Corp	Nanyang Technological University	Robert Bosch LLC
Argo AI	Center for Auto Safety	Intertek	NIO	UBER ATG
Aurora Innovations	Consumer Product Safety Commission	Liberty Mutual Insurance Company	Nissan North America Inc	UL LLC
AXA XL	Daimler Trucks North America	Locomation	Oak Ridge National Laboratory	University of York
Azevtec Inc	Edge Case Research	The MITRE Corp	Penn DoT	University of Waterloo
Babst, Calland, Clements & Zomnir	Infineon Technologies AG	Munich Re America	Renesas Electronics Europe GBMH	US DoT

Member organizations on UL 4600 Standards Technical Panel. Source: Underwriters Laboratories

WHITE PAPERS ON APPLICATIONS OF AI

Long-form content, produced by iMerit subject matter experts



CONTACT US



DATA ANNOTATION AND
AUTONOMOUS VEHICLE
SAFETY

imerit.net

