THE MITRE CORPORATION

SAFETY BUILDING BLOCKS OF HIGHLY AUTOMATED VEHICLES

by Zachary LaCelle and Dr. Christopher Hill

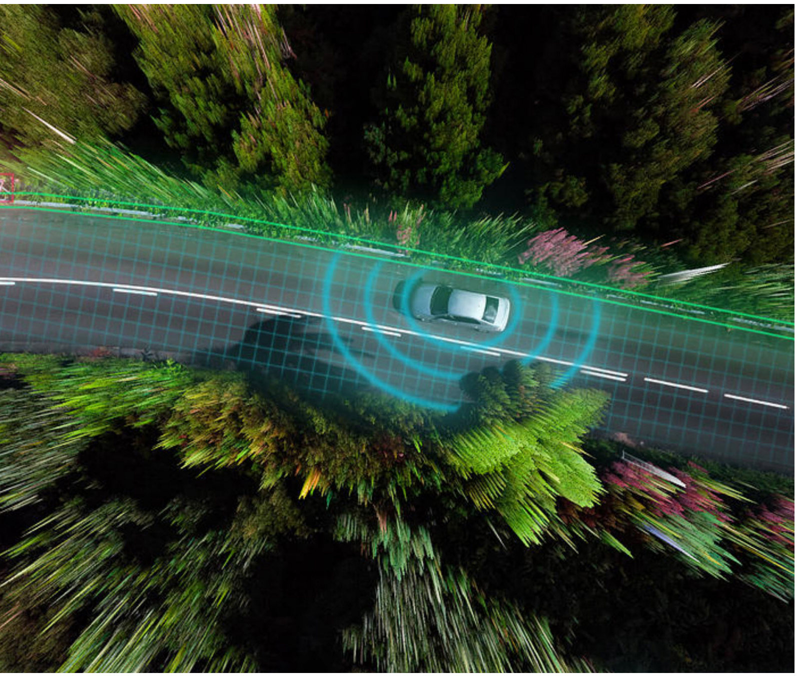

In ground vehicle transportation, the decade of the 2010s has shown the enormous potential of highly automated or autonomous vehicle technology. From initial automated systems such as adaptive cruise control and blind spot monitoring, innovators have expanded to the testing of fully autonomous systems that complete routes with no human input.[1]

As many researchers have highlighted, these technologies present opportunities in safety, accessibility, and efficiency for the entire transportation system.[2, 3] An entire new industry focused on these technologies has bloomed, with billions of dollars invested in non-traditional tech companies centered on autonomous vehicle technology.[4]

However, the technology required to accomplish these tasks is cutting edge and difficult to qualify as safe and market ready. In a system where humans expect highly reliable machines to keep them safe, these new innovations require advancements in evaluation methods. At lower levels of automation, such as Advanced Driving Assistance Systems like adaptive cruise control or lane-keeping assistance, human operators have had the task of ensuring safe performance of the system. However, research has shown that this method is flawed for more sophisticated systems; untrained drivers make poor safety operators.[5] As the technology evolves, relying on a human backstop for safety is insufficient. Furthermore, each year more systems capable of higher levels of autonomy, such as Tesla's Autopilot or Waymo One, are deployed—and further delay in policy approaches or regulatory frameworks means falling further behind the technology. A new paradigm is needed to ensure that the safety, accessibility, and efficiency gains promised by highly automated vehicles become a reality.

This new paradigm must be both flexible and holistic, recognizing that some fundamental challenges remain unanswered. However, now is the time to engage proactively and effectively to provide a clear and unambiguous set of recommendations, promising practices, requirements, and regulations around autonomous and automated driving systems (ADS). Indeed, recent actions such as the National Highway Traffic Safety Administration's (NHTSA) advanced notice of proposed rulemaking regarding safe adoption of ADS[6] underscore a sentiment throughout the industry: now is the time for a clear safety approach, cognizant of these unique challenges, that will remove environmental and regulatory uncertainty.

Challenges in ADS Deployment

Fully autonomous vehicles have not yet been successfully fielded at any large scale on roads today, despite many previous promises to the contrary. This is due simply to the scale of challenges facing these systems. Based on MITRE's research and prototyping experience—starting 15 years ago in the Defense Advanced Research Projects Agency (DARPA) Grand Challenge and continuing throughout the last decade with research in safety best practices, human-machine interfaces, trusted artificial intelligence, data-based hazard analysis for automated vehicles, and novel new approaches to autonomous perception, controls, and behaviors—MITRE finds that the following three challenges represent key roadblocks to safe and trusted ADS deployment.

The operational domain is incredibly complex: The roadway environment is highly cluttered, with many sizes and shapes of obstacles presented to drivers. These environments are also highly dynamic, with objects moving in and out of the roadway regularly. Thus, the sensing and perception challenges are significant. Unlike human drivers, who can classify things they have never seen before with relative accuracy, the current state-of-the-art systems used to detect and classify the environment do not yet adequately solve the problem for highly automated or autonomous systems. Furthermore, when the systems fail, they tend to fail unpredictably. Thus, the development and implementation challenges for safe perception systems remain fundamentally unsolved.

Human drivers cannot provide reliable failover for automated vehicles: Often, the approach taken by vendors is to assume that the human operator can serve as a backstop for the object and event detection and response task. This approach is used in Society of Automotive Engineers (SAE) level 3 automation, which allows for full automated driving but requires a human to monitor and take over when the ADS fails. Unfortunately, research has shown that humans make poor safety drivers.[7] The required time to obtain situational awareness,[5] combined with a driver's lack of formal training regarding their automated vehicle systems, means that systems relying on human fallback might actually be more dangerous than fully automated systems. It will be challenging for ADS to deploy without a comprehensive autonomy focused safety framework, since a staged deployment leveraging humans-in-the-loop may not be feasible.

"Miles driven" is insufficient to prove safety: To show system safety, often "miles driven" is the metric of choice. The thinking is that if an autonomous vehicle has operated with low rates of failure for thousands or millions of miles, surely it is safe to deploy on our roadways. Simulated miles, while very useful for autonomous vehicle development and testing, do not necessarily prove safety. The RAND Corporation has done research on the number of miles, without software and hardware changes, necessary to show that an autonomous system is safe—and those numbers are unachievably large.[8] Furthermore, due to the black-box nature of learning-based systems present on state-of-the-art autonomous vehicles, software is often updated during testing. After a software or decision model update, previous real or simulated miles may no longer indicate safety quality. Miles driven, while important, are not sufficient.

Therefore, due to the combination of an extremely challenging and dynamic environment, an inability to trust untrained human operators, and the insufficiency of miles driven metrics to prove safety, widescale deployment of these systems has remained elusive.

An Evolving ADS Safety Approach

The proposed new paradigm for ADS contains actionable recommendations for improved safety intelligence while recognizing that the approach is not yet complete. Autonomous vehicle systems have a very different set of strengths and weaknesses than human-operated systems, and future ADS safety approaches must be cognizant of these differences.

The proposed approach to ADS safety leverages key areas where these systems provide benefit over traditional vehicles. Specifically, **it leverages data-rich systems and fast-paced research and development to take an innovative approach to system safety—one that evolves and improves with new breakthroughs in safety research.**

Unfortunately, old methodologies—component-level and system-level functional safety combined with human-in-the-loop oversight—are not sufficient safety practices for ADS.[9] However, ADS present a wealth of opportunity in safety management through data-driven analysis and safety culture practices. These systems produce huge amounts of detailed operational data, far beyond that of a traditional vehicle. This data is also one of the only current methods to evaluate functionality and safety at scale; vehicle data is the only stand-in we have for the human operator's cognition. Thus, unlike traditional vehicle technology, a stronger and more prescriptive position must be taken regarding data logging, analysis, and sharing. Such an approach will serve as a key catalyst for systemic safety improvements and will encourage buy-in from all stakeholders such as researchers, regulators, and most importantly, the public.

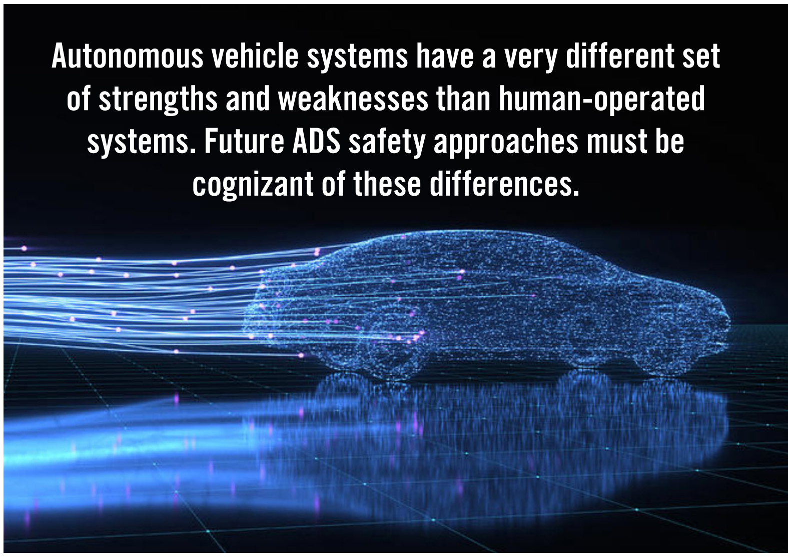

Autonomous vehicle systems have a very different set of strengths and weaknesses than human-operated systems. Future ADS safety approaches must be cognizant of these differences.

In addition to vehicle data, connectivity will play an important role in the ground vehicle fleet of the future. Stakeholders must work across government and with industry to enable connectivity wherever possible; since sensing the environment is challenging and often limited to line of sight, receiving shared communications about the environment greatly reduces risk and will thus increase deployment, adoption, and safety. This iterative and evolving approach is heavily focused on leveraging ADS data to inform safety, through methods such as domain-wide safety analytics or use-case sharing to demonstrate behaviors in common situations. By leveraging organizational structures that promote a safety culture and taking into consideration future ADS system requirements for communication and collaboration, MITRE believes that ADS development can move from mostly disparate and siloed efforts to a collaborative, vibrant, and safer ecosystem—providing benefit for all parties involved.

ADS Safety Building Blocks

Through MITRE's experience with safety systems in the aviation transportation domain, as well as our decades of research in autonomous and automated vehicle systems, the following initial building blocks for a safety framework have been identified. This list of building blocks is not yet complete: ADS safety remains an unsolved problem. However, these building blocks leverage known and proven safety approaches, as well as key ADS-enabling technologies, to bring down the safety risk industry-wide.

Safety Culture and Management

An organizational culture that proactively engages in safety risk management is critical to addressing and managing safety throughout design, development, and deployment of an ADS system. MITRE recommends that regulators collaborate with developers of ADS technologies to encourage organizational safety practices, such as the Safety Management System (SMS) approach, with possible consideration towards regulation. In December 2020, MITRE published "Management of Safety Risk in Automated Driving Systems",[10] which outlines how the SMS organizational approach can be applied to ADS development. This approach focuses on safety as a core cultural attribute of an organization, from the technical contributor up to company executives, and has been used in industries ranging from airlines to nuclear energy, and recently in the automotive industry.[11] ADS technology implementation is technically challenging, and often organizations focus solely on these technical hurdles—but safety must not be an afterthought. It must be considered from design all the way to deployment.

Assessing Safety through Data Sharing

Autonomous vehicles produce massive amounts of data—from sensor outputs, to complex models of the world, to control actions, to vehicle location information. This data provides insights into causes of hazards, both local to a specific ADS implementation and systemic across all vehicles. Because functional safety approaches for ADS cannot yet accurately measure or guarantee system safety, researchers, regulators, and developers must be able to assess performance and hazards effectively at scale. Data sharing partnerships help to address these challenges. Today, automobile manufacturers are already voluntarily collaborating on safety with each other and the NHTSA in a data-sharing partnership called the Partnership for Analytics Research in Traffic Safety (PARTS).[12] While currently focused on conducting analyses to gain insight into how advanced driver assistance systems (ADAS) perform in real-world scenarios, the PARTS vision is to expand to ADS. Vehicle content and other safety data is anonymized, pooled across organizations, and joined with police-reported crash information to provide data-driven safety insights—especially those related to system interdependencies. Data protection through a trusted third party builds trust from all participants, protecting driver privacy and preventing punitive responses to hazards which stifle safety reporting. Policymakers and regulators should work to ensure that PARTS successfully expands to address the needs of ADS and gains traction across the industry.

Hazard-Aware, Traceable Data Logging

Currently, there are no prescriptive requirements on data logging for autonomous vehicles. However, when an incident occurs, the only information available to identify root causes and mitigations is that data. In ADS systems, there is no guarantee of a human driver for safety investigators to question, whereas in non-fatal accidents involving traditional vehicles, interviews with drivers are a key component of the incident report. Manufacturers and ADS developers may be logging this information for various proprietary uses but, to fulfill the needs of safety analysts, this data must be tied to specific hazards, be traceable, and must be updated iteratively as new hazards are discovered. Current voluntary data-logging practices are insufficient; while event data recorders capture items such as vehicle acceleration or air bag deployments, they universally do not capture key ADS data elements—such as if the ADAS system is enabled. Thus, MITRE recommends that ADS developers be required to demonstrate a hazard-aware process for identifying and updating data elements within their data logger—to improve safety across ADS operations. All of this should be done while protecting consumer privacy.

Considering Communications, Spectrum, and Connected Vehicles

Due to challenges with initial deployment of connected vehicle communications, some ADS developers are designing without consideration for broad, cross-vendor connectivity. However, MITRE believes that connected vehicle technologies provide important safety and capability benefit. Therefore, to improve ADS safety, MITRE recommends continuing to push toward broadly deployed connected vehicle capabilities. To enable this ecosystem, regulators maintain dedicated spectrum resources for ADS technologies. The operating domain of ADS is incredibly challenging, and any capability that simplifies part of this domain is critical in safe deployment. As one example, identification and tracking of other motor vehicles requires a variety of cutting-edge sensing and perception technologies, as does detection of signage and other infrastructure. Vehicle to Everything connection technology (V2X) mitigates errors caused by failures in these technologies. Adoption and roll-out may be slow, as was seen with Dedicated Short-Range Communication technology; thus, aggressive protection of communications resources throughout the beginning of adoption is important. Additionally, if a new V2X implementation that is co-developed with industry and government partners shows promise in prototype tests, MITRE recommends that regulators promptly require the technology in vehicles that are SAE's automation level 3 or higher, as these systems execute complete vehicle control without effective low-latency human oversight.

Requiring Certification for Highly Automated Vehicles

Currently, regulatory oversight for ADS technology at the federal, state, and municipality levels has focused on guidance, recommendations, and best practices. As ADS technology increases in maturity and prevalence throughout the ground transportation space, MITRE expects that more prescriptive approaches will be required. NHTSA has recognized this in human-controlled vehicles with the Federal Motor Vehicle Safety Standards approach to codifying safety requirements for traditional automotive technologies. For ADS, the increased system complexity requires a delicate balance of requirements and regulations that do not stifle innovation and prevent technologists from making headway against the unsolved problem of autonomous driving. Thus, an evolving and flexible certification process should be used to provide a common set of requirements toward which to design. This is especially important for highly automated vehicles—systems that rely primarily on ADS for safety and functionality. Additional research and analysis are needed to determine what this certification process would entail, and would involve collaboration between industry leaders and regulators, with a focus on a performance-based and technology-agnostic process. As ADS systems are fielded, certification and independent review of ADS vehicles being deployed on our roads is necessary.

Conclusion

Automated driving technology will continue to proliferate across the transportation fleet. From traditional vehicle automation such as anti-lock braking, to lane keeping and adaptive cruise control, the driving population has shown a desire to simplify and improve their transportation experience. However, as control is shifted from human operators to computer algorithms, key safety considerations must be actively addressed. There remains no clear process to deploy new technologies, and the current waiver system will not scale. The approach and recommendations presented here offer a systematic method to improve ADS safety, recognizing the technological challenges that remain unsolved. By carefully and thoughtfully bringing evidence-driven sets of standards and practices to the challenges in ADS, with a focus on flexibility and adaptability in implementation, these autonomous vehicle systems can be safely and effectively deployed—realizing their safety and capability promises.

About the Authors

Zachary LaCelle is an Autonomous Systems Principal at The MITRE Corporation. He leads MITRE's Mobile Autonomous Systems Experimentation Lab, which serves as a resource for government and research partners enabling prototyping, experimentation, testing, and analysis of autonomous and automated systems—across the defense, security, and transportation domains.

Dr. Christopher Hill is the chief engineer of MITRE's Transportation Safety Division. He specializes in public-private partnerships and represents MITRE in the Partnership for Analytics Research in Traffic Safety (PARTS), a group dedicated to the advancement of traffic safety.

References

[1] I. Boudway, “Waymo Begins Fully Driverless Rides for All Arizona Customers,” 08 10 2020. [Online]. Available: https://www.bloomberg.com/news/articles/2020-10-08/waymo-one-app-offers-driverless-alternative-to-uber-in-arizona. [Accessed 31 12 2020].

[2] S. Kitajima, K. Shimono, J. Tajima, J. Antona-Makoshi and N. Uchida, “Multi-agent traffic simulations to estimate the impact of automated technologies on safety,” Traffic Injury Prevention, vol. 20, pp. S58-S64, 2019.

[3] Energetics, Inc. & Z, Inc, “Study of the Potential Energy Consumption Impacts of Connected and Automated Vehicles,” U.S. Energy Information Agency, 2017.

[4] D. Holland-Letz, M. Kasser, B. Kloss and T. Muller, “Start me up: where mobility investments are going,” McKinsey & Company, 2019.

[5] C. Gold, D. Dambock, L. Lorenz and K. Bengler, “Take over! How long does it take to get the driver back into the loop?,” Proceedings of the Human Factors and Ergonomics Society Annual Meeting, vol. 57, no. 1, pp. 1938-1942, 2013.

[6] National Highway Traffic Safety Administration, “Framework for Automated Driving Systems,” 19 11 2020. [Online]. Available: https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/ads_safety_principles_anprm_website_version.pdf. [Accessed 19 02 2021].

[7] P. Koopman and B. Osyk, “Safety Argument Considerations for Public Road Testing of Autonomous Vehicles,” SAE International Journal of Advances and Current Practices in Mobility, vol. 1, no. 2, pp. 512-523, 2019.

[8] N. Kalra and S. M. Paddock, “Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?,” RAND Corporation, 2016.

[9] P. Koopman and M. Wagner, “Challenges in Autonomous Vehicle Testing and Validation,” SAE International Journal of Transportation Safety, pp. 15-24, 2016.

[10] K. Hollinger and H. Shirazi, “Management of Safety Risk in Automated Driving Systems,” 2020. [Online]. Available: https://www.mitre.org/sites/default/files/publications/pr-20-3326-management-of-safety-risk-in-automated-driving-systems.pdf. [Accessed 29 1 2021].

[11] J. L. LaReau, “GM: We encourage employees, dealers to tattle after ignition switch crisis,” Detroit Free Press, 6 9 2019. [Online]. Available: https://www.freep.com/story/money/cars/general-motors/2019/09/06/gm-ignition-switch-nhtsa-recalls-safety-defects/2099289001/. [Accessed 29 1 2021].

[12] National Highway Traffic Safety Administration, “PARTS Partnership for Analytics Research in Traffic Safety,” [Online]. Available: https://www.nhtsa.gov/parts-partnership-for-analytics-research-in-traffic-safety. [Accessed 19 02 2021].

[13] ISO/IEC, “26262:2011,” Road Vehicles -- Functional Safety, 2011.

[14] Department of Defense, “Department of Defense Standard Practice, System Safety,” MIL-STD-882E, 2012.

[15] Underwriters Laboratories Inc., “ANSI/UL 4600”.Standard for Evaluation of Autonomous Products.

[16] N. Leveson, C. Fleming, M. Spencer and J. Thomas, “Safety Assessment of Complex, Software-Intensive Systems,” SAE International Journal of Aerospace, vol. 5, no. 1, pp. 233-244, 22 10 2012.

[17] M. Stoltz-Sundes, “STPA-Inspired Safety Analysis of Driver-Vehicle Interaction in Cooperative Driving Automation,” KTH Royal Institute of Technology, 2019.

[18] Waymo, “Wayo Safety Report,” 2020. [Online]. Available: https://storage.googleapis.com/sdc-prod/v1/safety-report/2020-09-waymo-safety-report.pdf. [Accessed 29 1 2021].

[19] National Transportation Safety Board, “Accident Report NTSB/HAR-17/02 PB2017-102600,” 2016. [Online]. Available: https://www.ntsb.gov/investigations/AccidentReports/Reports/HAR1702.pdf. [Accessed 29 1 2021].

[20] P. Koopman, U. Ferrell, F. Fratrik and M. Wagner, “A Safety Standard Approach for Fully Autonomous Vehicles,” WAISE, 2019.

[21] Federal Communications Commission, “In the Matter of Use of the 5.850-5.925 GHz Band: FIRST REPORT AND ORDER, FURTHER NOTICE OF PROPOSED RULEMAKING, AND ORDER OF PROPOSED MODIFICATION,” Washlington, D.C., 2020.

[22] W.K. Kellog Foundation, “Developing and Using a Logic Model,” Jan 2004. [Online]. Available: https://www.wkkf.org/resource-directory/resources/2004/01/logic-model-development-guide. [Accessed 01 Mar 2021].