A self driving license: Ensuring autonomous vehicles deliver on the promise of safer roads

Christopher Bradley^{1,2,†,*} and Victoria Preston^{1,2,†}

Edited by Serena Booth and Anthony Tabet

HIGHLIGHTS

- Existing regulations and standards for autonomous vehicles (AVs) fail to adequately consider highly autonomous systems.
- Current reliance on human drivers to be the “failsafe” for some AVs is problematic and requires reevaluation of current driver licensing procedures.
- AVs will only be safer than human drivers if appropriate support and testing protocols are provided and enforced.
- Thus, in order to enable safety certification for AVs, a multi-stage, data-driven oversight framework is proposed.

Upon maturation, autonomous vehicles (AVs) have the potential to provide significant benefit to society. A breadth of partially autonomous systems are already commercially available, and vehicles with advanced capabilities are tested and deployed on public roads. Although the advancement of AV technology is highly anticipated, the future of the industry currently rests on uncertain ground with respect to regulatory oversight. The current industry standards and legal regulations which apply to AVs are only equipped to fully ensure that simple autonomous capabilities are safe. As vehicles become more autonomous, and as driving decisions are shifted from human to computer, a regulatory paradigm shift will be necessary. This article reviews the present state of AV safety standards and regulations, and discusses the potential for regulatory evolution.

For decades, the prospect of autonomous vehicles (AVs) has attracted the attention of investors, government agencies, researchers, and consumers alike. While factors like convenience, accessibility, and aesthetic cannot be discounted as reasons for consumer interest in self-driving

vehicles, safety is without question the most immediately compelling incentive for widespread adoption [1]. In 2019 alone, approximately 4.4 million Americans were seriously injured in an automobile accident, with nearly 39,000 fatalities [2]. According to the United States Department of Transportation (DOT), from 2015 to 2016, approximately 94-96% of all serious motor vehicle crashes involved driver-related factors, such as impaired driving or illegal maneuvers [3]. Taking human error out of the equation with AVs has the potential to significantly reduce the number of serious accidents and save lives.

Although the idea of self-driving vehicles has existed nearly as long cars themselves, the race for autonomous driving shifted into high gear in 2004, when the Defense Advanced Research Projects Agency (DARPA) announced *The Grand Challenge* autonomous driving competition. Despite the fact that the initiative was motivated primarily by an interest in military applications for autonomous driving, it is credited for demonstrating that self-driving was not just an interesting theory, but a realizable objective. Sixteen years later, the commercial future of AVs looks bright. The modern AV industry is composed of dozens of corporations from around the world [4] worth an estimated \$54.23 billion USD in 2018 [5]. Several million vehicles with at least partially autonomous capabilities are predicted to be on the roads by 2024 [6], and by 2026 market projections suggest that AVs will be a \$556.67 billion USD global industry [5]. By 2030, advanced AV systems (i.e., fully driverless systems) are projected to be a \$60 billion USD industry alone, with nearly a third of the market dominated by North American corporations [7].

Initially, technological development for modern AVs enjoyed little to no regulatory oversight from federal or state governments. This laissez-faire approach to AV regulation was in part driven by a desire to hasten the arrival of a mature AV market, and has resulted in impressive progress in the field. In the span of a decade, adaptive cruise-control and emergency-braking have become nearly ubiquitous in new passenger cars. Even conditionally autonomous capabilities, in which the car performs the majority of certain driving tasks like parallel parking or lane changing, come standard in several mid-class consumer cars. Today, numerous companies are actively researching, testing, and deploying vehicles with advanced autonomous capabilities

¹Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA

²Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA

[†]These authors contributed equally.

*Email: cbrad@mit.edu

The authors declare no conflict of interest.

© 2020 The Author(s)

with the ultimate goal of full autonomy. However, the rapidly increasing number of on-road AVs, the changing landscape of autonomous software sophistication and capability, and the occurrence of several high-profile accidents and fatalities involving self-driving vehicles [8, 9], introduce some uncertainty for the future of the industry.

The ultimate impact of autonomous driving technology depends on reliability and public trust. Without either, the likelihood of widespread adoption of the technology declines, and the potential for the construction of legal roadblocks increases. For these reasons, the role of regulation and oversight for the industry must be carefully considered. Oversight frameworks can prevent overzealous companies from acting rashly and endangering lives, just as easily as they can create unnecessary burdens that stifle innovation. In this article, we discuss the current state of AV safety standards and regulations in the United States, and present several policy options for consideration by both legislative policymakers and industry leaders. In particular, we will focus our discussion on driver safety, software verification, and vehicle testing, all of which are necessary to prepare for the future filled with AVs.

Terminology: For the purposes of this article, the term *driver* will always refer to a human operator of a vehicle. *Passengers* are human occupants in a vehicle not involved in the driving task. When referring to a *vehicle* or *AV*, it could be in reference to any applicable ground transport system commercially made autonomous (e.g., passenger cars, buses, tractor-trailers).

Levels of Autonomy: The Society of Automotive Engineers (SAE) defines six levels of autonomy for AVs [10]:

- **Level 0 (L0): No Automation** A human performs all driving tasks at all times. Emergency braking or lane departure warnings are allowed.
- **Level 1 (L1): Driver Assistance** A driver controls the vehicle in the majority of driving tasks; the vehicle may feature a single automated systems (e.g., adaptive cruise control or lane-centering) to simplify operation.
- **Level 2 (L2): Partial Automation** A driver is monitoring and controlling driving tasks; the vehicle can perform multiple tasks at once (e.g., simultaneous adaptive cruise-control and lane-centering) as an advanced driver assistance system.
- **Level 3 (L3): Conditional Automation** The vehicle performs most driving tasks while the driver supervises. The driver is expected to handle specialized driving scenarios (e.g., unfamiliar surface streets) and be ready to intervene at all times.
- **Level 4 (L4): High Automation** The vehicle performs all driving tasks under specific circumstances (e.g., under certain speeds or weather conditions in pre-mapped regions) without need for direct supervision by the driver. Driver override is still available.
- **Level 5 (L5): Full Automation** The vehicle performs all

driving tasks under all circumstances, and no driver is necessary; passengers will never be involved in driving.

L0-L2 autonomy is nearly ubiquitous in modern cars, with features including adaptive cruise-control, emergency braking, and adaptive lane-assist. Within the last few years, vehicles with capabilities approaching those of L3, like Tesla's *Full Self-Driving* and GM's *Super Cruise*, have been available for purchase. Moreover, most industry research and development vehicles used on closed test-tracks and public roadways are rated L3+ [11, 12]. Even the testing of L4 vehicles without human oversight on public roads has begun in limited circumstances [13]–[15]. We next consider how these systems are currently being regulated, both internally by self-driving companies, and by federal, state, and local regulators.

AV Oversight Today

In the US, the authority to regulate AVs is shared by both the industry via internal voluntary standards, and the government at all levels (federal, state, and local) through both legislation and regulatory standards. Here we describe current standards and regulations, which are summarized in Table I.

Voluntary Standards: Design principles for automobiles of all kinds are borne primarily from standards and best practices defined by independent organizations. The International Organization for Standardization (ISO) is the preeminent authority on AV standards. ISO 26262 *Road Vehicles - Functional Safety* [16] provides best practice guidance and requirements for ensuring *functional safety* of a vehicle in the event of a system failure (e.g., electrical shorts). This regulation applies to systems like dynamic stability control, but does not cover hazard prevention or safety in the absence of a system failure, such as when a car automatically brakes. To better address non-failure safety cases presented by automated systems, ISO PAS 21448 *Safety of the Intended Functionality* (SOTIF) [17] was published in 2019. SOTIF aims to reduce the number of unaccounted for scenarios that automated systems may encounter. In addition to these core standards, a series of *Intelligent Transport Systems* standards from ISO are used to broadly define necessary attributes of cyberphysical systems (e.g., ISO 14825:2011 Geographic Data Files [18]). Together, the ISO standards are considered sufficient for up to L2 autonomous systems. In the US, independent standards relevant to autonomous vehicles are provided in principle by the SAE (e.g., SAE J 3016 which defines autonomy levels [10], or SAE J 3018 which outlines on-road testing procedures for L3+ vehicles [19]). Though these and other groups do not have explicit authority over vehicle manufacturers, they have considerable influence over how governments craft their own policies.

Federal: The US federal government has primarily published national standards and guidelines in an effort to influence companies, and form a basis for state and local policies. The Department of Defense (DOD) (e.g., MIL STD 882 Standard Practice: System Safety [20]) and the DOT National Highway Transportation Safety Administration (NHTSA under Title 49 of U.S. Code, Chapter 301 *Motor Vehicle Safety*)

each have published guidelines on, or which apply to, AV technology. NHSTA maintains and publishes *Federal Motor Vehicle Safety Standards* (FMVSS) [21] which largely hold vehicle manufacturers to a set of equipment requirements and certification standards. Relevant specifications for automated systems include Title 49: Chapter V, Part 571, Subpart B No. 124 Accelerator Control Systems, and Title 49: Chapter V, Part 571, Subpart B No. 121 Electronic Stability Control Systems. In 2017, NHSTA published AV 2.0 *A Vision for Safety* [22] which provided guidance on design, testing, and safety features on AVs and provided a voluntary framework for safety documentation. An update to the document, AV 3.0 *Preparing for the Future of Transportation* [3], was released in 2018 and provides guidance for training and licensing of test drivers. In January 2020, AV 4.0 *Ensuring American Leadership in Automated Vehicle Technologies* [23] was announced which expands on previous guidelines and primarily focuses on the role of the US Government in supporting technology growth. Although there are currently no legislative measures which apply specifically to AVs, House Bill H.R.3388 *Safely Ensuring Lives Future Deployment and Research in Vehicle Evolution Act* (or SELF DRIVE Act) [24] has been proposed (and approved by the House in 2017, currently under consideration in the Senate). The SELF DRIVE Act specifically aims to establish a set of federal regulation for AV safety on matters of cybersecurity, privacy, and accessibility, in addition to a system for safety assessment.

State: Traditionally in matters of automobile regulation in the United States, the federal government has held power over regulating physical components of a vehicle while states have determined licensing. Because of this, states currently wield most legislative influence over the AV industry, as approving driverless capabilities falls more naturally under the domain of licensing. The National Conference of State Legislators [25] reported that 40 states had either passed legislation, or issued executive orders, related to AVs at the time of writing this article. For example, in Massachusetts (MA), Part I, Title XIV, Chapter 90 of General Laws covers *Motor Vehicles and Aircraft*, and discusses registration standards, operational regulations (such as age or training of a driver or roadway rules), and sets regulations for minimum safety standards for functional components of a vehicle. State legislators may also pass restrictions for operation of an AV on motorways. In MA, Senate Bill S.2115 *An Act to promote the safe integration of autonomous vehicles into the transportation system of the Commonwealth* [26] is one such example of potential regulation of the use of AVs within public infrastructure. In addition to vehicle regulation, state governments are solely responsible for establishing criteria to earn driving licenses within the state (with the exception of commercial driving licenses which are federally administered). In states with companies testing AVs on roadways, additional oversight is imposed on certifying safety drivers for these companies (which can include additional background checks and proof of passing a company-designed curriculum for drivers).

Local: Finally, municipalities and other local governments play

a role in regulating or overseeing AV technology development on roadways. For example, the city of Boston released a vision and framework for development of AV technology *Autonomous Vehicles: Boston's Approach* related to an Executive Order from the Office of the Mayor *Establishing a Policy for Autonomous Vehicles in the City of Boston* [27]. The framework establishes a process by which Boston can oversee development and restrict access to the city's roads for self-driving companies. These efforts are coordinated at the state-level through the MA DOT in accordance to MA Executive Order 572 [28], and requires establishing a memorandum of agreement between the city and an AV company in order to operate. Local agreements between municipalities and self-driving companies are generally amicable, with benefits for both parties if the AV technology being developed is successful [29]. For some companies, connecting with city and state governments is a key part of their business model, with the intent of upgrading or supplementing existing public transit infrastructure.

The underlying message from federal guidelines, which influences state and local authorities, is that self-driving research and development should proceed with minimal restriction, while self-driving as a product for purchase should be regulated to the extent necessary to establish reasonable safety. State and local policies, by nature of being distributed, have an unfortunate consequence of potentially being at odds between regions. For example, states without any AV regulation (e.g., Maryland) may become more compelling places to launch new products than states with stringent regulations; or in an extreme example, states which allow L3+ commercial driving on highways may border states that do not (thus causing interrupted service for a driver). This consequence underscores the importance of unifying measures like the SELF DRIVE act.

Perhaps of more concern, however, is that the vast majority of regulations and standards discussed here are only sufficient for L1-L2 systems, despite the fact that L3+ systems exist on roads today. This regulatory gap is currently being addressed in two ways: by good faith "investment in safety" by AV companies, and through the development of American National Standards Institute (ANSI) and Underwriter Laboratories (UL) standard 4600 *Standard for Safety for the Evaluation of Autonomous Products* [30]. The former is best represented by *Safety First for Automated Driving* [31], which was published in 2019 by several AV companies and manufacturers. The document proposes a framework for testing, regulating, and certifying L3+ vehicles, and provides a model for allowing continued innovation while ensuring that capable and safe vehicles are available on the market. UL4600 is a proposed standard aimed specifically for L3+ technologies which places the burden of proof on companies to collect sufficient evidence that their offering is safe.

To proceed with L3+ technology development, placing the burden of proof on companies to demonstrate safety establishes the need for an *oversight* framework. In ideal conditions, oversight policies would encourage procedural

TABLE I: Independent standards and US regulations currently in place which apply to autonomous vehicles. This list is by no means exhaustive, but provides examples of standard or legislation at different scopes of influence.

Voluntary Standards	Standard Name	Description	Status	Citation
	ISO 26262	Functional Safety	Issued	[16]
	ISO PAS 21448	Intended Functionality (L1, L2)	Issued	[17]
	e.g., ISO 14825:2011, ISO 10711:2012	Intelligent Transport Infrastructure	Issued	[18, 32]
	SAE J3016	Autonomy Level Definitions	Issued	[10]
	SAE J3018	On-Road L3+ Testing	Issued	[19]
	MIL STD 882	US Transport System Safety	Issued	[20]
	ANSI/UL 4600	Evaluation of L3+ AVs	Proposed	[30]
	US DOT NHTSA FMVSS	US Min. Safety Requirements	Issued	[21]
	US DOT NHTSA AV 2.0, 3.0, 4.0	US AV Vision	Issued	[3, 22, 23]
	Safety First for Automated Driving	Informal Industry (US) Statement	Proposed	[31]
US Regulations	Regulation Name	Scope	Status	Citation
	House Bill H.R.3388 SELF DRIVE Act	Federal	Proposed	[24]
	e.g., MA Senate Bill S.2115	State	Proposed	[26]
	e.g., MA Executive Order 572	State	Issued	[28]
	e.g., Executive Order: Establishing a Policy for Autonomous Vehicles in the City of Boston	Local	Issued	[27]

clarity and transparency, and establish clear metrics for quantifying feature or behavior efficacy. One of the key challenges facing regulation and standardization of AVs today is the murkiness of what reasonable proof of safety looks like. This is particularly true for L3 systems, which are uniquely defined by collaboration between driver and vehicle for safe operation. Next, we highlight some of the challenges specific to L3 systems, then consider L4+ systems, and propose several oversight frameworks for consideration.

Human-in-the-Loop: Level 3 AVs

The day has not yet arrived when an AV can be summoned to your doorstep unattended and whisk you to any specified destination. In general, the vast majority of AVs on the road today are still controlled by a licensed driver as the primary vehicle operator, and any on-board autonomous system is typically designed to *augment* their capabilities. Systems like adaptive cruise control, emergency braking, lane correction, or evasive lane changing are designed to ease the mental and physical burden on the driver, particularly in safety critical scenarios in which the situational awareness of a driver may be compromised (e.g., merging car in a blindspot). *Situational awareness* describes the driver's mental model of everything in and around the vehicle, including the location and speeds of other cars, the position of their vehicle in a lane, and the road condition. L1 and L2 systems assist in cases where imperfect situational awareness can lead to accidents, acting in a supervisory role at all times. For L3+ vehicles, and for L3 AVs in particular, the operational responsibility of the vehicle shifts from the driver to the autonomous system, and the burden of safety-critical emergency response shifts to the driver. In this section, we highlight the key socio-technical challenges associated with L3 AVs, and suggest frameworks for regulation of L3 vehicles (and their drivers).

The Driver-Vehicle Relationship

As a "supervisor," a driver is removed from the direct operation of the vehicle unless performing a *takeover* operation, in which a safety critical or unknown scenario arises that the vehicle is explicitly not designed to handle. This requires the same level of situational awareness as driving, without any of the insight about the control trajectory of the vehicle the driver is in (since the driver is not the one making the decisions). In the L1-L2 scenario in which the vehicle is monitoring the driver, the vehicle is typically only monitoring a specific subset of safety scenarios, and does not evaluate *how* the driver is performing (e.g., if the driver is obeying all traffic rules, if the driver is taking the optimal trajectory around debris in the roadway, etc.). In the L3 scenario in which the driver is monitoring the autonomous system, the vehicle is capable of watching the blind spots and lane centering for itself just as before, but requires the driver to intervene if its actions are generally "unsafe" given the abstract context of a scenario or the rules of the road. It requires a different type of supervision from the driver. Furthermore, L3 assumes that the driver *will* need to intervene in some non-trivial number of scenarios. This relationship places a huge burden of responsibility on the driver, without giving the driver a sense of agency or control in their role.

Over the last 5 years, studies have examined the takeover experience for drivers of L3 AVs, and demonstrated that humans functioning in a supervisory role to an autonomous system require new skills or training to appropriately respond to takeover events. Autonomous monitoring systems, like those found in L2 vehicles, are explicitly designed to be consistent, vigilant, and dependable. In contrast, drivers possess unique characteristics that make their immediate reactions to takeover requests difficult to predict. For example, age [33] and emotional state [34] have been shown to be strong factors in the response of a driver to takeover

requests. Additionally, environmental complexity—such as traffic density [35]—has been shown to influence takeover time and intervention strategy adopted by drivers since situational awareness plays a critical role in selecting a safe action to take. Moreover, when drivers are allowed to engage in non-driving tasks, studies unilaterally suggest that vigilance is reduced and response time to takeover requests is slower than that of drivers engaged fully in the driving task [36, 37]. This does not take into account the case in which a driver may willfully ignore takeover requests that a vehicle makes, for example because they are confident that the vehicle can manage the scenario (perhaps based on past experience).

When companies test their L3+ vehicles on public roads, these vehicles are generally staffed by *safety drivers* who are specifically trained by the company on the operation of an AV. Much like typical driver education, the standard AV driving curriculum consists of both classroom instruction and on-road practice. Classroom instruction broadly goes over the interfaces in the vehicle, reviews takeover protocol and case studies, and highlights the importance of vigilance when monitoring the vehicle. On-road instruction typically consists of practicing safe takeover procedures. For the majority of AV companies, the process of selecting, training, and hiring safety drivers is a multi-week investment. One self-driving company, Aurora, details a six week course that every candidate must take in order to be certified by the company [38]. Other companies, like Waymo [39] and Uber [40] also provide extensive training for their safety drivers. Enrollment in such programs does not necessarily guarantee suitability of the individual as a safety driver; typically only a portion of applicants are considered suitable for full-time roles. The lengths that private companies go to train and test applicants are well supported by studies which show that drivers with experience performing takeovers through both in-classroom and simulated training are much more consistent and safe when responding to requests (e.g., [41]).

In many US states, these curricula are required to be put in place in order to allow operation of vehicles on public roads. For example, in MA, safety drivers must pass a background check performed by the state, must carry a valid driving license, provide proof of passing the company's safety driving curriculum, and have their status renewed by the company on a quarterly basis. Although the curriculum, and its metrics for "passing," are all set by the company, states have the ability to request changes to ensure public safety. The level of care and oversight for safety driving is well-founded: in nearly all AV accidents to date, driver error, distraction, or over-confidence in the autonomous system played a considerable role. Indeed, in a 2017 study [41] it was demonstrated that untrained drivers reported the highest confidence in an autonomy system both before and after they experienced a simulated safety critical takeover request. This raises the question about how L3+ technologies should be introduced to the broader public; indeed even the best L3 systems will require driver takeovers.

Mitigating Challenges Posed by L3 AVs

Arguably, the most dangerous part of L3 technology happens

at the interface of autonomous and human driving. In order to maintain continuous safe operation of the vehicle through these transitions, *driver engagement* must be maintained, and reliable *takeover hand-offs* need to be designed. As discussed, a primary cause of human failure in the role of a safety driver is caused by lapses in attention. To this end, AVs with the ability to monitor the driver for signs of attentiveness, or engage in a series of attention-keeping requests, have been proposed to improve driver readiness in the event of a safety-critical transition (companies Renovo Auto and Affectiva are two examples [42]). Interior facing cameras and biometric sensors which can identify if a driver has lost focus on the road can be useful tools in ensuring that the role of safety driver is being performed adequately.

Monitoring the driver is only one facet of creating safe takeover scenarios, however. When a takeover request is initiated, the driver needs to be able to recognize that the request is being made, understand how to engage with the vehicle in order to regain control, be able to infer what the reason for the takeover request was in order to act appropriately, and following the event, re-engage the control of the autonomy. A useful framework for understanding the relationship between the driver and the autonomous system is through the concept of *control authorities* which establishes the driver and vehicle as collaborators for a driving task, both of whom can initiate and execute actions [43, 44]. Interface standardization, regulated by NHTSA, is what makes it possible for a driver to operate any typical L0-L2 car without needing specific training. The accelerator is in the same place and does the same thing, whether the car is a Volkswagen manufactured 40 years ago, or a brand new Ford. In the same vein, alerts about potential takeovers and the way in which a takeover is initiated or ended should be required to follow the same principles. Walch et al. [45] extensively list many proposed interfaces to communicate takeover requests, and we point the reader to this source for further discussion. For drivers to be expected to operate an AV reliably, the physical interface for takeover requests must be standardized.

With this in mind, federal and state legislators must decide on if and how training for AV operation be included in driver licensing. Currently, the licensing system in each state largely focuses on training motorists on safety protocol while operating a vehicle: interpreting regulatory signs, understanding traffic laws, and assessing basic operational capability of the vehicle in multiple driving scenarios. Mastery of these skills alone does not imply a person will be an effective operator of a L3 AV. Possible pathways to AV driver training include augmenting current drivers education programs with AV classroom instruction or creating a specialized licensing path for AV owners/operators (in the style of commercial licensing, this could be federally standardized). The role of driver training could alternatively fall to the AV companies and dealers, as suggested by NHTSA AV 2.0 [22].

However, given the difficulty private companies have found training their own safety-drivers, it is worth questioning whether the average driver can be reasonably expected to

safely operate L3 AVs. Companies have been able to find individuals that are capable of performing this task, but there is informal consensus that the task is far more difficult than driving a non-autonomous vehicle. Even after rigorous vetting and training, human nature can lead safety drivers to over-trust the systems they are meant to monitor [46]. For this reason, some companies like Waymo, a leader in the current race for a fully autonomous vehicle, have committed to skipping commercial release of L3 vehicles entirely. There is no explicit reason why L3 vehicles must be adopted before L4+ vehicles, and solving the challenges unique to L3 technology (the driver-vehicle interface) do not lend insight to L4 capabilities. A unilateral ban on L3 vehicles for commercial purchase is yet another option worth serious consideration for regulatory measures. Leaders in safety of autonomous systems, like Dr. Mary Cummings of Duke University, support the idea that regulation is necessary for highly autonomous systems, and that L3 systems are potentially so dangerous that they should not be allowed on the roads [46]. Such a ban would not extend to test vehicles necessary for developing L4+ capabilities, so the progress of innovation of autonomous technology would not be seriously hindered. As L3 vehicles are already beginning to enter public roadways, consideration of this policy must be careful but swift.

Safe by Design: Certifying Level 4+ AVs

Though L3 vehicles present more immediate pressing policy questions, the long-term goal of self-driving technology is to expand into L4+. Such AVs are already being tested on public roads, and in some cases without a safety driver behind the wheel. Notably, in 2020 self-driving company Nuro was granted a FMVSS exemption for their vehicles, which permits the operation of vehicles without normally standard equipment, like rear-view mirrors [15]. Similarly, Waymo has tested its L4 vehicles without a safety driver in certain, carefully mapped, regions of Arizona [13, 14]. Fully driverless systems are a significant step beyond L3 systems in that by removing the need for a driver, engineers no longer need to be constrained by driver-vehicle interface requirements. In principle, the future of transportation could be one without a steering wheel at all. When cars are truly capable of driving themselves, regulatory focus must be capable of ensuring that the car will be capable of operating safely.

AV Certification is the process by which an AV is deemed “safe” according to standards and regulations imposed on that class of vehicle. L0-L2 vehicles today are certified via extensive crashworthiness testing, emissions tests, physical assessment of car components, and inspection for adherence to FMVSS and state regulations. Car companies are largely responsible for performing and reporting these tests to an oversight committee, which ultimately deems the vehicle suitable. An AV will likely be subject to many of these same tests, however, there are unique characteristics possessed by AVs which must additionally be considered. Next, we discuss these characteristics in addition to different frameworks for regulating AVs under the challenges posed. We also consider potential consequences of certain types of regulation.

AV Characteristics

Abstractly, an AV consists of a *body*, *perception system*, *datalogger*, *decision-making systems*, and *control system*. The body is the physical frame and wheels of the vehicle, and can include the interior (seats), windshield, and mirrors. AVs may not look like typical vehicles; for example, there is no explicit need for a windshield if cameras are being used to “see” the road. The perception system encompasses the sensors (and their computers) used to collect data about the environment the vehicle is in. Typical sensor suites include LIDAR, RADAR, and cameras. The perception system fuses measurements from different sensors together into a model of the world. This model is served to the decision-making system which is a computer (or multiple computers) which plans trajectories and actions the vehicle should take according to a navigation goal and the model of the environment. Trajectory plans are provided to the control system which actuates the motor, brakes, and steering for the vehicle to physically execute these plans. The datalogger is used to record raw sensor information, messages passed between the different computers, and decisions made by the vehicle. Together, these physical components form a network of multiple computers which process data in order to create appropriate actuator commands.

Testing these components for faults, and establishing safe behaviors in the event of such a fault, is largely covered by standards set in ISO 26262. This includes establishing redundant components and pathways for data transfer and electrical power. However, an AV could “work” perfectly, and still engage in unsafe behavior. These non-fault scenarios are what ISO/PAS 21448 examines for L1 and L2 systems, however, it does not provide guidance suitable for AVs in which the majority of driving responsibility falls on the AV and not the driver. For instance, there is significantly more complexity in the process used to transform a LIDAR return into a steering command. Generally, the lack of standards and regulations for fully self-driving AVs is reflective of the relative infancy of technology in this space. Despite this, it is inevitable that a system for certifying AVs will become necessary, and is central in the argument for the development of ANSI/UL4600.

Further complicating the design of a certification scheme is the capability and prevalence of *over-the-air (OTA) updates*, in which changes to the software on an AV can be made remotely and automatically, and which can change the behavior of the AV without any changes to the hardware¹. Engineers are constantly working on improvements to the sensing, decision-making, and control systems onboard a vehicle. With more vehicle miles, comes more data that can be used to improve generalizability of the AV to more scenarios. However, OTA updates generally break traditional models of vehicle certification. The same car may

¹It is worth noting that cybersecurity is an incredibly important topic as it relates to AV safety. The same systems that enable OTA updates make AVs vulnerable to attack. Cybersecurity is out of scope for this article, but we direct the interested reader to the NHTSA landing page to get started <https://www.nhtsa.gov/technology-innovation/vehicle-cybersecurity>.

be updated to behave so differently from prior to the update, that it may need to be re-certified. How updates can be incrementally certified in order to continuously improve AV technology is currently an open question. Based on our current understanding of the field today, several potential frameworks for assessing AV technology, which consider all of the characteristics and the challenges associated with them, have emerged.

Frameworks for Ensuring AV Safety

Algorithms used for decision-making are difficult to inspect; hundreds of thousands of lines of code may be used to define behaviors or models of an AV. Further complicating matters is the common use of machine learning techniques in self-driving. A model of the vehicle's situational awareness or decision-making process may be "learned" from data and experience, but the exact way that decisions are made is obfuscated (known as a "black box"). Thus, it is impossible to know with certainty the behavior of a vehicle in any arbitrary situation. Because of this, thorough and diverse testing is necessary to define the probability of safe operation of the vehicle. Here, we describe four potential frameworks to assess safety of L4+ AVs. Fig. 1 summarizes these approaches.

a) *Scenario Testing, A Driving Test for AVs:* One approach to self-driving certification would be to create a standardized "driving test" for AVs to complete in order to be approved for public roads. This approach, which may rely on both simulated and real driving elements, could be an easy way to filter out under-performing AVs on the basis

Figure 1: Four frameworks for AV safety oversight are proposed, clockwise from top left: Certifying Components (CC), Holistic Certification (HC), Scenario Testing (ST), and Graduated Capabilities (GC). In CC, specific behaviors of an AV, like navigation or detection, are inspected. In contrast, HC uses global metrics like "Driving Infractions per Mile" to quantify the vehicle's general efficacy. ST examines an AV's behavior in specific scenarios like driving on empty streets or crossing intersections. Finally, GC combines different aspects of the other frameworks in a "capabilities" based framework, in which an abstract capability like "localization" could be tested through a series of increasingly complex simulated and field trials. Images in this figure are from the KITTI raw dataset [47].

of clearly established expectations on specific tasks. The International Telecommunications Union (ITU) Focus Group on AI for Autonomous and Assisted Driving in the European Union (EU) announced in Fall 2019 its plan to assemble an AV driving test for highly automated vehicles [48]. While this approach may be appealing for its familiarity and simplicity, there are several reasons it can not be a stand-alone test of autonomous vehicles.

One of the challenges of defining a good driving test for AVs is fundamental to a key limitation of autonomous systems: transferability. A licensed driver can be trusted after training and a single on-road test on a rainy day in rural Massachusetts, to reasonably handle their vehicle on a sunny day in San Francisco, without too much concern. However, an AV that uses deep learning systems, or any other data-driven approach, cannot necessarily be expected to transfer its driving skills between conditions or scenarios that it has never seen before. Perception, and the deep networks they rely on, can be especially susceptible to deviations from data in a training set [49]. If a perception system is trained on a set of scenarios that do not fully represent the space of all realistic scenarios, it is not safe to deploy, even though it may pass a "self-driving test." Demonstrating the occasional fragility of computer vision, research has shown that detection of stop signs can be made unreliable by simply placing a few strips of tape over the road-signs [50]. Although progress has been made in the area of learning robust features for computer vision [51], generalizability remains a major roadblock to certification of autonomous systems. An autonomous driving test would need to be augmented by further verification in order to properly assess the capabilities of an AV in a reasonable diversity of scenarios.

b) *Certifying Components:* Another approach to AV safety testing would be to perform a subsystems level review, that would propose to individually certify that each separate subsystem of the AV operates as expected. This would mean, for example, that the perception system alone would be tested in detection tasks for cars, trees, signs, or pedestrians, and the accuracy scored. Perhaps the greatest advantage to this methodology is that incremental updates which only impact some subsystems would only require re-testing those subsystems. Potentially billions of hours of on-road data would not be needed to certify the vehicle.

Unfortunately, this approach alone leaves a glaring hole in safety testing: subsystems on an AV interact, and at these interfaces there is opportunity for insidious failures. An infamous example of a failure that would not be caught in this framework is exemplified by a crash with an Uber vehicle in Tempe, AZ [8]. In this case, the perception system failed to classify a pedestrian consistently, and, while the "object" was tracked the whole time, the velocity trajectory of the pedestrian was erased for each new classification label (it is important to note that though a safety-driver was present, they did not intervene due to external distraction). This is an example of the perception system and the decision-making and control systems creating a compounding unsafe behavior

in their interaction. Since it is impossible to guarantee that all systems will be perfect, the cascading impact of error or uncertainty *will* be a major factor for AV safety².

c) *Holistic Certification*: The faults of the component-based approach naturally lead to the third possible framework for certifying an L4+ vehicle: holistic certification. Unlike scenario testing, which assesses the AV behavior on a subset of specific tasks, holistic certification represents a more general approach to proving the safety of an AV through data-driven metrics. Possible metrics could include collision frequency or driving law infractions per mile.

In aviation, another highly-automated field, one metric is *chance of failure per hour*, and sets the standard at a maximum of 10^{-9} failures per hour. A similar standard for AVs could be put into place. To meet the threshold set for aviation would take 10^9 hours of driving, which is equivalent to 30 billion miles of driving for an average car. In one study, it was shown that hundreds of millions of miles of testing would be necessary to confidently state fatality, injury, and crash rates. Notably, for high-confidence estimates of performing *better* than a human, 11 billion, 161 million, and 65 million miles would be necessary for fatality, injury, and crashes respectively [54]. For reference, the total number of miles driven by Waymo's fleet since 2009 as of January 2020 was 20 million miles on roadways and is substantially short of these proposals [39].

As in aviation, simulation is necessary to reach the prescribed test hours for AVs. While simulation is a useful tool, simulation can differ from the real world in meaningful ways. Visual appearance, underlying physics, and the actions of other simulated vehicles may not be fully captured in simulation. In addition, simulation scenarios would need to be defined to test a breadth of autonomous behaviors, and studies would need to be completed to determine how well the simulation performance generalizes to the real world performance of the vehicle. One method for developing such simulations is to use data collected in the field perturbed by noise or slightly modified scenarios. An alternative to standard "photorealistic" simulation, which attempts to replace on-road testing, is to perform hardware-in-the-loop fault simulations, which stochastically introduce problems to the autonomous software at a much higher rate than might occur naturally. These approaches are similar to fault injection simulators (e.g., [55]) or software stress testing (e.g., [56]), and could reduce the burden of proof about the safety of autonomous systems to an attainable level.

d) *Graduated Capability Certification*: A fourth approach could be to establish procedures for testing the

vehicle holistically on *capabilities* of the vehicle. This is a natural combination of component, scenario, and holistic testing frameworks. This approach focuses specifically on the idea of defining abstract capabilities for an AV to possess. One such capability is localization: the ability for the vehicle to be accurately located (within a map). This approach is proposed in *Safety First for Automated Driving* [31] and uses ISO/PAS 21448 to structure a 4-part graduated certification process: (1) analysis, (2) verification, (3) validation, and (4) field operation. In this framework to certify that an AV can localize itself, a study of the sensor error models would be conducted (analysis), failsafe behaviors would be tested for sensor dropout or malfunction (verification), simulation scenarios would be evaluated (validation), and several on-road tests would be performed (field operation). Central to using this framework is creating a list of capabilities an AV should have, and establishing criteria in each of the 4 parts to improve the capability and prove, within reason, that the vehicle possess that capability.

Realistic Oversight

The suggested frameworks are complementary ways in which AV safety can be considered. Each of these methods relies on significant data collection from on-road and simulation testing, and places the burden of proof on AV companies/manufacturers to defend their systems. The diversity of approaches AV companies are taking to design autonomous systems, in addition to the myriad use cases for AVs (e.g., public transit in geo-fenced areas, highway travel), makes an external, independent review process or strict regulation unrealistic. Just as ANSI/UL4600 proponents suggest, it is likely through standard-setting and external oversight that safety can be ensured for AVs of all levels of autonomy and functions.

Collaborative efforts between standards groups, legislators, and companies have the potential for appropriately supporting technical innovation while protecting public safety, however, there are some failure cases that should be highlighted. The first relates to setting standards with impossible criteria to meet. As it stands, it is not possible to prove with certainty that an AV will "always" or "never" do something. Any legislation with such language must be re-written to account for the stochastic nature of any real-world system. A second relates to lobbying guidelines, which will inherently favor one company's approach over others. Some significant fear in the industry is that larger companies with more clout could potentially box out competition by setting standards that only the most mature technologies in the field can meet. To avoid either of these scenarios, an effective oversight system is one that incrementally establishes guidelines as the field as a whole matures. To this end, the direction that the field moves can be set by industry standards, and legal regulations can be put in place when a reasonable minimum threshold for safety is proven by most companies. Precedents for these guidelines can be assessed by external review boards which receive data and analysis from company testing procedures.

²Risk-bounded behavior design and analysis of uncertainty propagation in autonomous systems could be used to mitigate the effects of compounding error. *Uncertainty propagation* quantifies the impact of process or measurement uncertainty downstream in a decision-making or modeling systems. *Risk-bounded* planning encompasses systems that explicitly characterize and consider uncertainty in the estimation of an autonomous agent's state (e.g., [52, 53]). Quantifying and using uncertainty is one of the key technical challenges that AV developers tackle.

Discussion

For non-autonomous vehicle regulation, international and federal standards dominate. This ensures that every vehicle on the road today meets a minimum threshold of safety to protect the general public. Drivers are trained by the states, but with the implicit understanding that drivers trained in one state will be able to safely operate their vehicle in another. As driving tasks are offloaded to the vehicle in self-driving systems, the separation between certifying vehicle safety and certifying safe vehicle operation no longer applies. The new landscape for vehicle regulation and standardization will require considerably more oversight from international and federal bodies. House Bill H.R.3388 the SELF DRIVE act, has been one recent attempt at creating such a system. The relative infancy of self-driving, however, makes large regulatory changes, like those proposed in H.R.3388 difficult to pass. ANSI/UL4600 is another such national-level attempt at setting standards for highly autonomous systems which has yet to be fully approved and adopted at the date of publication.

In this article, we described a path toward incremental regulatory measures by first discussing necessary near-term measures for L3 AV technology, particularly with respect to preparing drivers for a paradigm shift in the way human-vehicle interactions will occur on the road. We further discussed multiple frameworks for regulating the autonomous decision-making systems of L4+ AVs, and highlighted the unique regulatory challenges posed by software-intensive systems. In particular, we suggested that, as a whole, verifying the safety of AVs will require careful collaboration between standards agencies, governments, and companies in order to appropriately identify performance criteria to define safe-operation of an AV. This collaboration is practically necessary given the complexity and opaqueness of autonomous algorithms and models, which are fundamentally impossible to characterize with certainty.

Beyond the topics raised in this article, it is worth highlighting that the future of self-driving is tied not only to the success of verifiable software and safe driver training, but also to factors less directly related to the vehicles themselves. Urban planning, for example, can simplify the problem of self-driving by making roads friendlier for AVs. Additionally, cybersecurity is a major issue with AVs, especially those which update and change their software frequently. The safety of the vehicle depends on the fact that the software being tested is not vulnerable to outside attack. There is also a long-standing question about *liability* with regards to AVs (i.e., the person or entity to blame in the case of an accident) that could influence safe practice adoption at companies and manufacturers. In short, autonomous driving has the potential to completely change the literal landscape in our cities, and figuratively with respect to regulation and insurance.

The arrival of fully autonomous vehicles on our roads is generally accepted to be a matter of when, not if. Although L5 AVs are not anticipated until sometime in the next several decades, L3 and L4 systems are already on the roads. For an industry with hundreds of billions of dollars of invested

capital, and where the prevailing opinion is that the first to develop a fully autonomous vehicle will have a massive market advantage, companies have significant incentive to accelerate innovation. At the same time, the safety of the general public must be considered and protected. In principle, these two goals are complementary; a successful self-driving vehicle should be one that is safer than the car of today. Regulation and standardization for AV technology is necessary not to limit innovation, but to set the goalposts for suitably capable vehicles. The success of AV technology will be overwhelmingly tied to strategic standardization and regulation of the industry.

Citation

Bradley, C. & Preston, V. A self driving license: Ensuring autonomous vehicles deliver on the promise of safer roads. *MIT Science Policy Review* 1, 76-86 (2020).

Open Access

This *MIT Science Policy Review* article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- [1] Early adopters of self-driving cars motivated by time, traffic hassles - torcrobotics. URL <https://torc.ai/early-adopters-of-self-driving-cars-motivated-by-time-traffic-hassle>. (Accessed on 06/27/2020).
- [2] National Highway Traffic Safety Administration, Department of Transportation. National Center for Statistics and Analysis (NCSA) motor vehicle traffic crash data (2020). data retrieved from CrashStats, <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812946>. Accessed on 05/13/2020.
- [3] Department of Transportation, National Highway Traffic Safety Administration. Automated vehicles 3.0 preparing for the future of transportation (2018). URL <https://www.transportation.gov/sites/dot.gov/files/docs/policy-initiatives/automated-vehicles/320711/preparing-future-transportation-automated-vehicle-30.pdf>. Last accessed on 05/13/2020.
- [4] CB Insights. 40+ corporations working on autonomous vehicles (2020). URL <https://www.cbinsights.com/research/autonomous-driverless-vehicles-corporations-list/>. Last accessed on 05/14/2020.
- [5] Jadhav, A. Autonomous vehicle market by level of automation (level 3, level 4, and level 5) and component (hardware, software, and service) and application (civil, robo taxi, self-driving bus, ride share, self-driving truck, and ride hail) -

- global opportunity analysis and industry forecast, 2019-2026 (2018). URL <https://www.alliedmarketresearch.com/autonomous-vehicle-market>. Last accessed on 05/14/2020.
- [6] Semi-autonomous and autonomous vehicles market by level of automation (level 1, level 2&3 autonomous vehicles), ADAS features, components (AI, radar, LIDAR, camera, and ultrasonic sensor), fuel type (ICE, EV), and region - global forecast to 2022 & 2030 (2017). URL <https://www.marketsandmarkets.com/Market-Reports/near-autonomous-passenger-car-market-1220.html>. Last accessed on 05/14/2020.
- [7] Wagner, I. Projected size of the global autonomous car market in 2018 and 2030 (2020). URL <https://www.statista.com/statistics/428692/projected-size-of-global-autonomous-vehicle-market-by-vehicle-type/>. Last accessed on 05/14/2020.
- [8] Causland, P. Self-driving Uber car that hit and killed woman did not recognize that pedestrians jaywalk. *NBC News* (2019). URL <https://www.nbcnews.com/tech/tech-news/self-driving-uber-car-hit-killed-woman-did-not-recognize-n1079281>. Last accessed on 05/14/2020.
- [9] List of self-driving car fatalities - Wikipedia. https://en.wikipedia.org/wiki/List_of_self-driving_car_fatalities. (Accessed on 07/03/2020).
- [10] SAE-J3016_201806. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. Standard, SAE International (2018).
- [11] Autopilot and full self-driving capability | Tesla. <https://www.tesla.com/support/autopilot>. (Accessed on 07/03/2020).
- [12] Super cruise - hands free driving | Cadillac ownership. <https://www.cadillac.com/world-of-cadillac/innovation/super-cruise>. (Accessed on 07/03/2020).
- [13] Markman, J. Real autonomous cars hit the road in Arizona. *Forbes* (2019). URL <https://www.forbes.com/sites/jonmarkman/2019/11/23/real-autonomous-cars-hit-the-road-in-arizona/#53a7cc3132c6>. Last accessed on 07/03/2020.
- [14] Executive order 2018-04: Advancing autonomous vehicle testing and operating; prioritizing public safety (2018). URL <https://apps.azdot.gov/files/sitefinity-files/Executive-Order-2018-04.pdf>. (Last accessed on 03/15/2020).
- [15] Hawkins, A. J. The federal government just granted its first driverless car exemption. *The Verge* (2020). URL <https://www.theverge.com/2020/2/6/21125358/nuro-self-driving-delivery-robot-r2-fmvss-exemption>. Last accessed on 05/13/2020.
- [16] ISO-26262-1:2018. Road Vehicles — Functional Safety. Standard, International Organization for Standardization, Geneva, CH (2018).
- [17] ISO/PAS-21448:2019. Road Vehicles — Safety of the Intended Functionality. Standard, International Organization for Standardization, Geneva, CH (2019).
- [18] ISO-14825:2011. Intelligent Transport Systems — Geographic Data Files (GDF) — GDF5.0. Standard, International Organization for Standardization, Geneva, CH (2011).
- [19] SAE-J3018_201909. Safety-Relevant Guidance for On-Road Testing of SAE Levels 3, 4, and 5 Prototype Automated Driving System (ADS)-Operated Vehicles. Standard, SAE International (2019).
- [20] MIL-STD-882E. Standard Practice for System Safety. Standard, US Department of Defense (2012).
- [21] FMVSS. Federal Motor Vehicle Safety Standards and Regulations. Standard, US Department of Transportation, National Highway Traffic Safety Administration (2010).
- [22] Department of Transportation, National Highway Traffic Safety Administration. Automated driving systems 2.0: A vision for safety (2017). URL https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/13069a-ads2.0_090617_v9a_tag.pdf. Last Accessed on 05/13/2020.
- [23] Department of Transportation, National Science and Technology Council. Automated vehicles 4.0 ensuring american leadership in automated vehicle technologies (2020). URL <https://www.transportation.gov/sites/dot.gov/files/2020-02/EnsuringAmericanLeadershipAVTech4.pdf>. Last accessed on 05/15/2020.
- [24] SELF DRIVE Act, H.R.3388, 115th Cong. (2017).
- [25] Autonomous vehicles state bill tracking database (2020). URL <https://www.ncsl.org/research/transportation/autonomous-vehicles-legislative-database.aspx>. (Last accessed on 05/08/2020).
- [26] 191st General Court of the Commonwealth of Massachusetts. S.2115, An act to promote the safe integration of autonomous vehicles into the transportation system of the commonwealth (2019). URL <https://malegislature.gov/Bills/191/S2115>.
- [27] Office of the Mayor Martin J. Walsh. Executive order: Establishing a policy for autonomous vehicles in the city of Boston (2016). URL https://www.boston.gov/sites/default/files/embed/file/2016-10/executive_order_autonomous_cars.pdf. Last accessed 08/04/2020.
- [28] Governor Charles Baker. No. 572: To promote the testing and development of highly automated driving technologies (2016). URL <https://www.mass.gov/executive-orders/no-572-to-promote-the-testing-and-deployment-of-highly-automated-driving>. Last accessed on 05/15/2020.
- [29] Testing automated driving systems in massachusetts. URL <https://www.mass.gov/guides/testing-automated-driving-systems-in-massachusetts>. Last accessed on 05/13/2020.
- [30] ANSI/UL-4600. Standard for Evaluation of Autonomous Products. Standard, Underwriters Laboratories (2020).
- [31] Wood, M. *et al.* Safety first for automated driving (2019). URL <https://www.aprive.com/docs/default-source/white-papers/safety-first-for-automated-driving-aptive-white-paper.pdf>. Last accessed on 05/14/2020.
- [32] ISO-10711:2012. Intelligent Transport Systems — Interface Protocol and Message Set Definition between Traffic Signal Controllers and Detectors. Standard, International Organization for Standardization, Geneva, CH (2012).
- [33] Clark, H. & Feng, J. Age differences in the takeover of vehicle control and engagement in non-driving-related activities in simulated driving with conditional automation. *Accident Analysis & Prevention* **106**, 468–479 (2017). <https://doi.org/10.1016/j.aap.2016.08.027>.
- [34] Du, N. *et al.* Examining the effects of emotional valence and arousal on takeover performance in conditionally automated driving. *Transportation Research Part C: Emerging Technologies* **112**, 78–87 (2020). <https://doi.org/10.1016/j.trc.2020.01.006>.
- [35] Gold, C., Körber, M., Lechner, D. & Bengler, K. Taking over control from highly automated vehicles in complex traffic situations: the role of traffic density. *Human Factors* **58**, 642–652 (2016). <https://doi.org/10.1177/0018720816634226>.
- [36] Eriksson, A. & Stanton, N. A. Takeover time in highly automated vehicles: noncritical transitions to and from manual control. *Human Factors* **59**, 689–705 (2017). <https://doi.org/10.1177/0018720816685832>.
- [37] Petersen, L., Robert, L., Yang, J. & Tilbury, D. Situational awareness, driver's trust in automated driving systems and secondary task performance. *SAE International Journal of Connected and Autonomous Vehicles* **2**, 129–141 (2019). <https://doi.org/10.4271/12-02-02-0009>.
- [38] Aurora. Vssa_final.pdf. URL <https://downloads.ctfassets.net/>

- v3j0gnq3qxwi/4QVMTwpBo2ZOmE03B09UP2/611de2c139aef05d7204ace06e946e00/VSSA_Final.pdf. Last accessed on 05/14/2020.
- [39] Waymo. Waymo safety report: On the road to fully self-driving (2018). URL <https://waymo.com/safety/>. Last accessed on 05/15/2020.
- [40] Uber Advanced Technologies Group. A principled approach to safety (2018). URL <https://docs.huihoo.com/car/Uber-ATGSafety-Report-2018.pdf>. Last accessed on 05/15/2020.
- [41] Hergeth, S., Lorenz, L. & Krems, J. F. Prior familiarization with takeover requests affects drivers' takeover performance and automation trust. *Human Factors* **59**, 457–470 (2017). <https://doi.org/10.1177/0018720816678714>.
- [42] Hawkins, A. J. We need self-driving cars that can monitor our attention along with our emotions. *The Verge* (2018). URL <https://www.theverge.com/2018/4/13/17233614/renovo-affectiva-self-driving-car-driver-monitoring-emotion>. Last accessed on 05/13/2020.
- [43] Lu, Z. & de Winter, J. C. A review and framework of control authority transitions in automated driving. *Procedia Manufacturing* **3**, 2510–2517 (2015). <https://doi.org/10.1016/j.promfg.2015.07.513>.
- [44] Flemisch, F. *et al.* Towards a dynamic balance between humans and automation: authority, ability, responsibility and control in shared and cooperative control situations. *Cognition, Technology & Work* **14**, 3–18 (2012). <https://doi.org/10.1007/s10111-011-0191-6>.
- [45] Walch, M. *et al.* From car-driver-handovers to cooperative interfaces: Visions for driver–vehicle interaction in automated driving. In *Automotive User Interfaces*, 273–294 (Springer, 2017). https://doi.org/10.1007/978-3-319-49448-7_10.
- [46] Korosec, K., Roy, A. & Niedermeyer, E. W. Autonocast the future of transportation: #83: Dr. Missy Cummings of Duke University. *The Autonocast* (2018). URL <http://www.autonocast.com/blog/2018/6/12/83-dr-missy-cummings-of-duke-university>.
- [47] Geiger, A., Lenz, P., Stiller, C. & Urtasun, R. Vision meets robotics: The KITTI dataset. *International Journal of Robotics Research* **32**, 1231–1237 (2013). <https://doi.org/10.1177/0278364913491297>.
- [48] International Telecommunication Union. Focus group on AI for autonomous and assisted driving (FG-AI4AD) (2019). URL <https://www.itu.int/en/ITU-T/focusgroups/ai4ad/Pages/default.aspx>. Last accessed on 06/26/2020.
- [49] Azulay, A. & Weiss, Y. Why do deep convolutional networks generalize so poorly to small image transformations? *arXiv preprint arXiv:1805.12177* (2018). <https://arxiv.org/pdf/1805.12177.pdf>.
- [50] Hacking street signs with stickers could confuse self-driving cars | Ars Technica. <https://arstechnica.com/cars/2017/09/hacking-street-signs-with-stickers-could-confuse-self-driving-cars/>. (Accessed on 06/30/2020).
- [51] Madry, A., Makelov, A., Schmidt, L., Tsipras, D. & Vladu, A. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations* (2018).
- [52] Ono, M., Williams, B. C. & Blackmore, L. Probabilistic planning for continuous dynamic systems under bounded risk. *Journal of Artificial Intelligence Research* **46**, 511–577 (2013). <https://doi.org/10.1613/jair.3893>.
- [53] Guibas, L. J., Hsu, D., Kurniawati, H. & Rehman, E. Bounded uncertainty roadmaps for path planning. In *Algorithmic Foundation of Robotics VIII*, 199–215 (Springer, 2009). https://doi.org/10.1007/978-3-642-00312-7_13.
- [54] Kalra, N. & Paddock, S. M. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transportation Research Part A: Policy and Practice* **94**, 182–193 (2016). https://www.rand.org/pubs/research_reports/RR1478.html.
- [55] Jha, S., Banerjee, S. S., Cyriac, J., Kalbarczyk, Z. T. & Iyer, R. K. AVFI: Fault injection for autonomous vehicles. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*, 55–56 (IEEE, 2018). <https://ieeexplore.ieee.org/document/8416212>.
- [56] Hutchison, C. *et al.* Robustness testing of autonomy software. In *IEEE/ACM International Conference on Software Engineering: Software Engineering in Practice Track*, 276–285 (IEEE, 2018). <https://doi.org/10.1145/3183519.3183534>.