Response to Request for Comments (Federal Register No. 2020-02332)

Ensuring American Leadership in Automated Vehicle Technologies: Automated Vehicles 4.0

Steven E. Shladover, Sc.D.

(Personal perspective of a retired annuitant from the University of California PATH Program)

From the perspective of one who has been working on automated vehicle technologies since 1973, it is encouraging to see that so many agencies of the federal government are paying attention to these technologies now. This subject has been in and out of fashion in several cycles during my professional career, but the current level of interest is certainly higher than any that I can recall. This level of interest is of course encouraging on one hand, but it also brings some serious risks on the other hand, because wishful thinking about the potential future benefits of Automated Driving Systems (ADS) has outstripped realistic recognition of the severe technical challenges that remain to be overcome before ADS can be deployed safely.

Several important issues appear to have been overlooked in the development of *Automated Vehicles 4.0* and the federal government actions to date:

(1) Immaturity of the ADS technology requires a skeptical view of claims about its capabilities
(2) ADS do not fit well within the existing motor vehicle safety regulatory paradigm because they combine both vehicle and driver functions
(3) Many ADS developers coming from the IT industry, in addition to the traditional automotive industry, do not have a safety culture and are susceptible to pressures to "cut corners" on safety
(4) There is therefore a serious risk to public safety from under-regulation of ADS.

1. Technological Immaturity

The DOT policy documents regarding ADS have contained many statements that indicate, implicitly or explicitly, that DOT policy makers believe that the ADS technology has progressed so rapidly that it is nearly ready for widespread deployment. Although great progress has been made on the technology within the past decade, based on tens of billions of dollars of industry investments, several fundamental technological challenges have still not been overcome. Even the most advanced and sophisticated ADS developers, building on their billions of dollars' worth of effort, still cannot demonstrate safety comparable to average human drivers, much less achieving the level of skilled and attentive drivers. The only publicly available data that can provide an indication of the maturity of the ADS technology are the California disengagement reports. These show that even the handful of most advanced systems currently under test are requiring their test drivers to disengage the ADS at average intervals of once per 10,000 to 20,000 miles of driving to avoid safety-critical events (which have not been described precisely by the system developers). The large majority of the companies developing ADS systems have far more frequent disengagements than this, indicating that they are well behind the industry leaders. To put these numbers in context for comparison with today's human drivers, property damage-only crashes occur about once for every 300,000 miles, injury crashes about once for every 2 million miles and fatal crashes about once for every 100 million miles of driving. Based on just this elementary comparison, policy decisions must not be founded on *a priori* assumptions that ADS are safer than human drivers. Comparisons of ADS safety with human driver safety must be based on solid evidence and peer-reviewed analyses rather than wishful thinking, especially when they are serving as the basis for regulatory policy decisions.

The ADS developers have been successful in learning how to handle the most common scenarios encountered in driving on public roads, but they still have much work to do to manage the more uncommon scenarios that account for the most dangerous traffic situations. While the progress that the leading ADS developers made on improving their disengagement statistics was rapid several years ago, the rate of improvement has decreased in the last couple of years, as they have had to grapple with the remaining increasingly difficult and rare scenarios. We should expect this trend to continue because of the inherent difficulty of approaching the safety record of average human drivers. Even the ADS developers with the least frequent disengagements still require safety-related disengagements at a rate more than one order of magnitude higher than the rate of occurrence of property damage crashes (fender benders).

It is important to distinguish between the technological reality and the reports that appear in the media (both general interest and trade media) and especially those on the Internet. The ADS developers and the industry financial analysts have been hyping their descriptions of the state of development and of their readiness for widespread deployment for close to a decade in order to stimulate more investment and pump up the values of their companies. We can see this with the benefit of year 2020 hindsight, because during the 2012-2013 ramp-up of interest in ADS, companies were making dramatic predictions about what would happen by 2020. Now that we are in 2020 it should be obvious that those predictions were far removed from the reality of where we stand today. Why should we assign much credibility to the predictions that the same people and companies are making today, when we can see how far off their previous predictions were? Indeed, within the past year, we have finally started to see some of the largest and most advanced companies within the ADS industry adjusting their predictions to be much more conservative than they were a few years ago. It appears to take about a decade of hard work and some billions of dollars of investment for organizations to gain sufficient experience to recognize how much they still don't know. Those who have not yet committed this big a stake are the ones who are more likely to be making over-hyped claims today.

ADS developers will need to earn the trust of the general public and their representatives in the consumer and traffic safety interest groups by providing solid enough data about the safety of their systems to withstand critical scrutiny by technical experts. This means that the ADS developers will need to show how they have properly accounted in their design processes for the hazards that their systems will encounter in real-world driving and how they have tested their systems to demonstrate safety that is no less than the safety of a reasonably skilled and attentive human driver. That will require sharing of considerable data, including data that may be embarrassing when it indicates some problems or limitations, and it may also be sensitive in terms of the intellectual property of the developers. A delicate balance needs to be found to permit sharing of enough information to convince justifiably skeptical stakeholders that the system is safe, without disclosing so much information that it could jeopardize the competitive position of the developer. This is an area in which federal regulations are likely to be needed, to provide a level playing field for all industry participants and a solid fact-based foundation of knowledge upon which to base well-informed public decisions about ADS safety. As in the case of new drugs, convincing evidence of the safety of each ADS should be required before the public is exposed to them in regular use (without the supervision of a safety driver).

2. Inadequate Current Motor Vehicle Safety Regulatory Paradigm

The existing U.S. motor vehicle safety regulatory paradigm of federal regulations covering vehicle design and construction and state regulations covering driver performance and qualifications does not work for ADS, in which the driving behavior is governed by software embedded in devices built into the vehicles. The traditional boundary between vehicle and driver is violated by ADS, so a different regulatory paradigm is going to be needed.

In *Automated Driving Systems 2.0*, NHTSA asserted its primacy over regulations for ADS: "Allowing NHTSA **alone** to regulate the safety design and **performance aspects of ADS technology** will help avoid conflicting Federal and State laws and regulations that could impede deployment" (p. 18, emphases added). It also sought to deter states from even codifying its own guidelines to ADS developers: "NHTSA strongly encourages States not to codify this Voluntary Guidance (that is, incorporate it into State statutes) as a legal requirement for any phases of development, testing, or deployment of ADSs." In *Automated Vehicles 3.0*, the Department of Transportation expressed strong reluctance to creation of regulations at the federal level: "U.S. DOT … will reserve nonprescriptive, performance-based regulations for when they are necessary." (p. 41), indicating that they do not believe that regulations are necessary now (and leaving unresolved the question of how a regulation could be "nonprescriptive"). Nothing in the current *Automated Vehicles 4.0* document indicates that this approach has changed, which appears to leave DOT opposed to the creation of any regulations at any level of government that would address the safety of ADS. *Automated Vehicles 4.0* includes general statements about NHTSA's role in creating and enforcing safety standards (pp. 8 and 22) but no statements indicating an intention to create standards for ADS safety.

It is bad enough that DOT does not want to do its job of creating safety regulations to govern the most important recent innovations in motor vehicle technology, but the current federal policy would also leave state and local governments unable to defend their citizens from the hazards that would be created by immature ADS that have not been fully vetted for safety. State and local governments can currently decide to place reduced speed limits on sections of roadway and to restrict HazMat vehicles from driving on some parts of their road networks, and of course the states also decide which drivers should be licensed to drive on their roads based on criteria that they define. Comparable authority over driving safety would not appear to apply to vehicles driven by ADS if the federal government were to pre-empt state and local authority to regulate them, while failing to create significant new federal regulations to fill the regulatory void surrounding safe driving behavior of ADS.

The existing federal "self-certification" mechanism for motor vehicle safety regulations is based on the vehicle developers certifying that they comply with specific FMVSS requirements and meet their test criteria. This is relatively straightforward for traditional vehicle mechanical and electrical technologies and the FMVSS that have been defined to cover them, but no FMVSS have been created (or are even on the horizon) to cover ADS driving behaviors. This means that there is no basis upon which ADS developers can "self certify" the safety of their systems. There are not even any industry consensus standards that could be referenced for self-certification of ADS compliance. Although Appendix C to the *Automated Vehicles 3.0* policy statement contained a long list of industry standards, NONE of the published standards that were listed there apply to the safety-related performance of ADS. A few of the ongoing standards activities listed in Appendix C will eventually lead to relevant standards for narrowly defined ADS use cases, but even after those are published they will not come close to being a meaningful basis for self certification that could provide assurance of safety.

The "self certification" mechanism does not work in the absence of standards with specific requirements that the certifier can claim to satisfy. Specifically, what could the ADS developer actually certify? Would they simply certify that their ADS is "safe"? This would be analogous to the 16-year-old novice driver certifying to their local DMV that they are a safe driver and therefore demanding a driving license. States use varying combinations of written tests and on-road driving tests to judge whether novice drivers are ready to be given a driver's license. Unfortunately, there is no valid analog to this that can be applied to ADS. There have long been hopes for development of a driving test that could be applied to ADS to determine that they are ready for public use. For a variety of technical reasons, this is not yet viable and I believe it would not be worth investing the considerable effort that would be needed to try to develop such a test within the foreseeable future. Each ADS will be designed for a different operational design domain (ODD), and any such test would need to be tailored to the relevant ODD for the intended target use. More seriously, in order to be a safe driver, the ADS will need to successfully handle a very wide range of hazardous situations, and it would be extremely costly and time consuming to design and execute a test that could cover enough of those to produce convincing evidence of safety. Once a test was designed, ADS developers would be motivated to "design to the test", so the ADS that pass all of the test conditions could still be unsafe if they cannot also handle the wide range of hazardous conditions that are not included in the test.

Given this situation, I think that the best approach to authorizing ADS for public use within the immediate future is to focus on providing assurance that the ADS developer has followed a proper safety-conscious development process. This means that the developers should be required by federal regulations to comply with the existing safety development process standards that are most applicable to ADS, specifically ISO 26262, ISO PAS 21448 (Safety of the Intended Functionality – SOTIF), and UL 4600. This can be done within the existing self-certification framework by requiring the ADS developers to self-certify that they have followed the processes defined in all of those standards. If they have not followed some of the provisions in these standards (which is very likely, given the complexity of following them all), they should be required to submit a solid written technical explanation for why they have not followed those specific provisions, which can then be assessed by the regulator in deciding whether their process still provides adequate assurance of safety.

3. Safety Culture Shortcomings

The language of *Automated Vehicles 4.0* and of its immediate predecessor federal government policy statements about "Automated Vehicles" (or more precisely, ADS) is focused on the potential safety (and other) benefits that could be gained from deployment of the systems, without evident acknowledgment of the potential safety risks from deployment of systems that have not been thoroughly engineered. Responsible regulatory agencies need to be even more focused on the "worst case" scenarios than on the "best case" scenarios. The implicit assumptions that all the developers of ADS are not only highly technically skilled but also extremely conscientious about prioritizing the safety of their products need to be reconsidered from a more critical perspective.

Although the traditional automotive industry (OEMs and their suppliers) have long experience in developing safety-critical systems and testing them thoroughly before releasing their products, they are not the only active participants in ADS development. Many new entrants from the information technology world have become very active in ADS, especially from Silicon Valley, where there has been minimal experience with safety-critical systems. The core expertise of most of these companies is in computer science, software development, "artificial intelligence" and related fields, but not in the integration of complex electro-mechanical systems, human factors or system safety. The large majority of their technical staff members lack experience in developing systems that could kill or injure people when they don't work correctly. Key aspects of Silicon Valley culture are indeed antithetical to the development of safety-critical systems:

- Emphasis on speed of development rather than thoroughness (exemplified by the popular mantra, "move fast and break things");
- Releasing products before they have been thoroughly tested so that customers serve as the beta testers to find faults;
- Investors expecting very rapid and large returns on their investments (and demonstrating their impatience by threatening to "pull the plug" on funding if products are not released by prescribed deadlines);
- Rapid turnover of technical staff;
- Ego-Maniac CEOs making extravagant claims to investors and media to pump up their stock prices and then pressuring staff to try to match those claims, even when this is not technically feasible;

- Hyper-competitive attitudes to market timing, so that if a competitor has announced a product release they need to accelerate their own product release, even if the product is not really ready.

This culture has produced remarkable progress in fields like mobile phone applications, internet search, social media and enterprise software solutions, but those do not have the safety-of-life criticality of ADS. The attitudes and development processes that have succeeded in those domains are likely to be disastrous if applied to ADS, leading to systems that diminish, rather than enhance, the safety of driving.

If Automated Driving Systems are intended to lead to significant improvements in traffic safety, we should not be looking to Silicon Valley for guidance but should rather be thinking about the lessons that can be learned from the transportation modes that have already demonstrated significantly better safety than highway traffic – rail and air transportation. The air transportation example is particularly instructive because it was indeed quite unsafe in the mid-20th century, when both industry and government recognized that the future of the industry would be in jeopardy unless its safety record was improved dramatically.

The safety of the air transportation system was improved dramatically through close coordination between government and industry, with a very strong regulatory push from the FAA. There was a recognition that public infrastructure was needed to coordinate the use of the airspace to avoid mid-air collisions and to enable the network and terminal areas to function safely and efficiently (the air traffic control system). Strict requirements for training, certification and re-certification of vehicle operators and for maintenance and operations of vehicles were established under FAA rules. Aircraft were required to go through rigorous government certification procedures before the FAA would issue airworthiness certificates. The recent unfortunate experiences with the Boeing 737 MAX have shown what can happen when regulatory rigor is relaxed, industry prioritizes speed of deployment and cost reduction over safety, and regulators place too much trust in the industry.

It would be difficult and costly to scale the FAA regulatory framework directly to the highway transportation system, but this still provides an important data point regarding how to achieve large safety improvements in a complex transportation system. It would be naïve in the extreme to assume that such safety improvements will just happen through "voluntary" industry action in the absence of regulatory pressure. The creation of NHTSA in 1966 was in large part a reaction to the revelations about the unsafe designs of the popular passenger cars of the day, and the implementation of FMVSS since then have led to significant improvements in the safety of vehicle mechanical and electrical systems. A new generation of regulations will be needed to ensure the safety of the new generation of software-based automated driving systems, whose behavior is not covered by any current regulations.

4. The Risk of Under-Regulation is More Severe than the Risk Over-Regulation

ADS, if designed and deployed with due attention to their safety criticality, have the potential to improve traffic safety. However, if designed and deployed without due attention to their safety criticality, they have the potential to create new traffic safety hazards. Some of the ADS developers will give safety the attention it demands and some of them will not. It is naïve to assume otherwise.

Meaningful and sensible regulations should not impose significant burdens on the ADS developers who are already doing a good job of safety management. However, those regulations are necessary to modify the behavior of the ADS developers who are not doing a good job of ensuring the safety of their systems, whether that is because of technical or managerial shortcomings. When ADS developers put immature and unsafe systems on the road, they endanger not only the occupants of their vehicles but also all the other road users (vehicle occupants, pedestrians, cyclists) in their vicinity. This will have a secondary multiplier effect that can be even more damaging in the long term. The Uber and Tesla crashes have already raised significant public concerns about ADS safety, with good justification. Just a few more of these kinds of events could turn public attitudes hostile toward ADS, preventing or at least significantly delaying future potential safety benefits. The unsafe actions of the worst "bad actors" in the industry can destroy public confidence in the entire industry even before their safety incidents reach the frequency that would trigger a NHTSA defects investigation or recall – it would be too late by then.

The federal government appears to want to pre-empt state and local regulatory constraints on ADS to make it easier for developers to deploy the same systems everywhere, but at the same time the recent AV policy statements have shown an unwillingness to implement federal regulations. The purely voluntary approach emphasized in *AV2.0* and *AV3.0* is totally inadequate to constrain the behaviors of the companies that lack the technical, managerial or ethical qualifications to implement safe systems. The federal government needs to set significant minimum safety requirements and also allow states or locals to demand even better safety to provide additional protection to their citizens based on more challenging local conditions or local attitudes toward managing risk.

The greatest risk to public safety at this time in the ADS domain is a regulatory "race to the bottom" that caters to industry preferences to avoid regulations. This will ultimately be damaging to the industry as well as to public safety. The states are already under pressure to be lenient on regulations to attract what they expect to be significant new technology jobs, without always understanding the safety risks to their citizens. Those states that recognize the risks need to have the flexibility to provide greater protection to their citizens, while the citizens of those states whose governments are blind to the risks deserve adequate regulatory protection from the federal government.