TOYOTA

Statement of

Gill Pratt

CEO

Toyota Research Institute

on

"Self-Driving Cars: Road to Deployment"

before the

U.S. House of Representatives

Committee on Energy & Commerce

Subcommittee on Digital Commerce
and Consumer Protection

February 14, 2017

Chairman Latta, Ranking Member Schakowsky, and members of the subcommittee, thank you for the opportunity to appear before you today on this important topic.

My name is Gill Pratt, and I am the CEO of the Toyota Research Institute (TRI). Before joining Toyota, I was a professor at MIT, a founder of Franklin W. Olin College of Engineering, and a Program Manager at DARPA (the U.S. Defense Advanced Research Projects Agency), where I led the DARPA Robotics Challenge. As you may be aware, prior to my tenure at DARPA, in 2004, 2005, and 2007, the DARPA Grand and Urban Challenges developed fundamental technology to show the world that autonomous driving was feasible.

TRI was formed in January of 2016, and is focused on the development and advancement of artificial intelligence and related technologies. At TRI, we have four goals: (1) to greatly enhance vehicle safety and someday create a car that is incapable of causing a crash; (2) to greatly increase mobility access for those who cannot drive; (3) to develop robots to improve quality of life, particularly for older people; and (4) to accelerate discovery in materials science. TRI is located wholly within the United States, with its headquarters in Palo Alto, California, and additional teams in Ann Arbor, Michigan, and Cambridge, Massachusetts. Our initial 5-year budget commitment from Toyota is $1 billion. In addition to our own work, we sponsor approximately $20 million of research every year at U.S. universities, including Stanford, the University of Michigan, MIT, and many others.

In order to achieve our goals, TRI is intensely focused on the development of autonomous vehicle technology. TRI's focus builds on Toyota's long history with autonomous technology. In fact, Toyota has been generating vehicle patents in the autonomy field in the U.S. since 2006.

According to a report last year by the Intellectual Property and Science division of Thomson Reuters, Toyota is – far and away – the global leader in the number of self-driving car patents.

Autonomous vehicle technology is expected to significantly improve vehicle safety. In 2015 alone, approximately 1.25 million people died globally in automobile crashes, including 35,092 people in the U.S. This means that about the same number of people who died on 9/11 die every day around the world in car crashes. Many, many more are injured. Because more than 90% of crashes are caused by human error, autonomous vehicle technology has the potential to dramatically reduce these numbers. In addition, a decline in traffic-related deaths and injuries is likely to have significant economic benefits, such as reduced medical costs and less time and productivity wasted in traffic jams. By one estimate, the traffic congestion that could be alleviated by autonomous driving technology would save the U.S. economy $124 billion, or $1700 per American household. Reduced traffic congestion will also have environmental benefits by increasing fuel economy and reducing greenhouse gas emissions.

We also anticipate that autonomous vehicle technology will have a profound positive impact on older people and people with disabilities. Providing mobility options to those who cannot drive is likely to deliver important social and economic benefits. These include fostering independence, providing opportunities to participate in community and social activities, and reducing barriers to employment.

Two Paths to Autonomy

TRI is currently pursuing two paths to autonomy – a system called Guardian and a system called Chauffeur. Under Guardian, the autonomous vehicle technology acts as an always-watching crash mitigation system. Guardian operates in the background and is constantly monitoring the environment, stepping in when it perceives a collision is imminent. Simple examples of the Guardian approach include automatic emergency braking, which is standard equipment in almost every Toyota model that will be sold in the U.S. this year. Automatic emergency braking attempts to prevent or mitigate frontal crashes by applying the brakes if a frontal collision is imminent.

Under Chauffeur, the autonomous technology takes over the driving task from the human driver. Because the technology – rather than a human driver – is driving the vehicle, Chauffeur is an important step in achieving mobility for people who cannot currently drive.

Much of the hardware and software being developed for Guardian and Chauffeur is the same. The difference is that Guardian only engages when needed, while Chauffeur is engaged at all times during autonomous driving. Nevertheless, we believe there is an important role for both systems. For example, the Guardian approach enables us to introduce higher levels of driver assistance into our production vehicles in the near-term, helping to save more lives sooner, as we continue our progress towards Chauffeur.

We are currently testing and refining both Guardian and Chauffeur. To date, most of our testing of Chauffeur has been done on closed courses in a number of states. We are testing on public roads in Michigan, and have plans to test on public roads in California and Massachusetts after we receive regulatory approval from the states to do so. Because these systems can save lives, our hope is to deploy our systems as soon as possible, but we will only do so when we know that they can be deployed safely and responsibly.

Common Vocabulary

At present, there is quite a bit of confusion among the media, policymakers, and the public about autonomous vehicle technology. Before embarking on a discussion of the opportunities and challenges associated with testing and deploying these systems, it is important to take a few moments to make sure that we share a common understanding and vocabulary.

Under the SAE International Recommended Practice J3016, there are five demarcated levels of autonomy. All automakers are aiming to achieve Level 5, where the system can drive under any traffic or weather condition in any place and at any time. Although this is a wonderful goal, none of us in the automotive or information technology industries are close to achieving Level 5. Current prototype vehicles can handle many situations, but there are many other scenarios that are simply beyond current machine competence. It will take many years of research and development and many more miles of machine learning and testing to achieve the performance and reliability required for Level 5.

Level 4 is less capable than Level 5 because the autonomy only operates in specific situations. This may include limited areas of operation, limited speeds, limited times of day, or limited weather conditions. When companies say that they hope or intend to deploy fully autonomous vehicles in the next few years, they are typically referring to Level 4.

Level 3 is similar to Level 4, but the autonomy must sometimes hand off control to a human driver. Hand-off is a difficult challenge because the human driver may be engaged in other tasks and not paying attention. As defined by the SAE, in Level 3, the autonomy must give the driver sufficient warning of the need for a hand-off and must detect any condition requiring a hand-off. Because both of these requirements are extremely difficult to guarantee, it is possible that Level 3 may be as difficult to accomplish as higher levels of autonomy.

In Level 2, a vehicle hand-off to a human driver may occur at any time with only a second or two of warning. This means that the human driver must be able to react, mentally and physically, at a moment's notice. Moreover, a Level 2 system does not guarantee that it will always detect when a disengagement is necessary, so the driver must remain vigilant and monitor the road ahead – even when the autonomy is engaged. For example, a Level 2 system may fail to recognize and react to certain types of debris that fall from a vehicle traveling in front of it. It would be the responsibility of the human driver to not only notice the falling debris, but take over operation of the vehicle from the autonomy system – all in a split second. As you may be aware, some Level 2 systems have already been put in consumer vehicles.

Level 1 encompasses the driver assistance features we see in many vehicles today, such as adaptive cruise control, parking assistance with automatic steering, and lane keeping assistance. Under Level 1, most driving functions are still controlled by a human driver, but a specific function (such as steering or accelerating) can be done by the vehicle. Level 1 systems can be very sophisticated, such as our Guardian concept, or very simple, such as cruise control. In Level 1, the human driver is always engaged in the driving task.

As policy is developed to govern autonomous vehicle technology, it is important that all stakeholders share a common vocabulary. We appreciate that the National Highway Traffic Safety Administration (NHTSA) adopted the SAE's definitions in its recent Federal Automated Vehicle Policy. To ensure consistency and clarity, we urge all legislators and regulators – at both the state and federal level – to use the SAE definitions and taxonomy in their policymaking efforts relating to both the testing and deployment of autonomous vehicle technology.

How Safe is Safe Enough?

Society tolerates a significant amount of human error on our roads. We are, after all, only human. On the other hand, we expect machines to perform much better. Humans show nearly zero tolerance for injuries or deaths caused by flaws in a machine. However, the artificial intelligence systems on which autonomous vehicle technology will depend are presently and unavoidably imperfect. So, the question is “how safe is safe enough” for this technology to be deployed.

As I mentioned previously, there were more than 35,000 fatalities on U.S. roads involving vehicles controlled by human drivers. What if we could create an autonomous vehicle system that was as safe, on average, as human drivers? Would that be safe enough, particularly if it resulted in greater convenience, less traffic, and less impact on the environment? Because we judge machines more critically than we judge each other, the answer is probably no.

What if the machine was twice as safe as human-driven cars and only 17,500 lives were lost in the U.S. every year? Would we accept such autonomy? The answer is probably still no.

Policymakers – working with industry and relevant stakeholders – must determine what constitutes a sufficient level of safety for autonomous vehicle technology. As we sit here today, it is not clear how this measure will be devised or by whom. Perhaps as important, it is not currently clear whether it will be consistent across the entire U.S., let alone the entire world. However, before developers can complete testing of these systems and deploy the technology, this foundational question will need to be answered.

Technology Development and Testing

Policymakers must keep in mind that testing is a necessary means to an end. The goal for all developers of this technology, including TRI, is to develop a vehicle that can be deployed safely and responsibly. We cannot reach that goal unless we are able to test our technology in real-world environments, including on public roads. Testing is what will allow us to determine when our technology achieves a sufficient level of performance and is ready for deployment. Without making public roads available for extensive testing, we risk companies or entities deploying autonomous vehicle technology that is not yet ready for prime time.

Policymakers cannot expect that a system in development will perform at the same level as a finished system that is ready for deployment, but should expect a developer to take reasonable steps to protect the public when a system in development fails or does not achieve required performance. This could include requirements for a trained safety driver or operator who can monitor and - if necessary - take over operation of the vehicle being tested, as well as the ability to achieve a minimal risk condition or failsafe mode should the driver or operator be unable to take over operation.

It is important to remember that, for decades, automakers have been allowed to test advanced safety technologies on public roads. This includes the testing of sophisticated automated technology, such as adaptive cruise control and automatic emergency braking. While we appreciate that autonomous vehicles are of greater interest to the media and the public, we do not believe that it is necessary to entirely transform the process that governs how automakers test their new safety technologies. We recognize that autonomous vehicle technology is being developed by companies and entities other than traditional automakers that lack the same track record of safely testing on public roads. We also understand that policymakers and regulators want confidence that all testing will be conducted responsibly and safely. The challenge for policymakers is to enable companies that have demonstrated themselves to be responsible developers and testers of automotive safety technology, while also creating appropriate safeguards with respect to those that are less experienced.

State Patchwork

One of the most significant challenges that we face today with respect to the testing of autonomous vehicle technology is the patchwork of policy initiatives at the state level. More and more states are developing legislation and regulations that are unfortunately creating impediments to the development of autonomous vehicle technology.

We appreciate the good intentions behind most of these efforts. We understand that many policymakers are trying to spur innovation in their state or to prevent their state from being left behind as the technology flourishes elsewhere. However, these legislative proposals are likely to have the unintended and opposite result of discouraging development or investment in favor of a state with a less restrictive or more permissive regulatory framework. For example, as I noted previously, we are currently testing our systems on public roads in Michigan – which has implemented a very supportive regulatory framework – but have not yet initiated public road testing in California or Massachusetts – which have both implemented a more restrictive regulatory framework.

A number of proposed state regulatory frameworks veer into territory that has traditionally been the purview of the federal government – namely, vehicle safety performance standards. Traditionally, the driver has been the responsibility of the states and the vehicle has been the responsibility of the federal government. We recognize that the simplicity of this traditional dividing line is challenged as vehicles become more automated and the vehicle itself becomes the driver. We also understand that, without clear or certain direction from the federal government, some states may wish to take action to regulate the safety of these systems. However, we firmly believe that the establishment of vehicle performance standards for autonomous vehicle technology should take place at the national level.

Encouraging and incentivizing manufacturers to test autonomous vehicle technology in a wide variety of environments should be a primary objective of policymakers concerned with the safety of these systems. Driving in Silicon Valley is not the same as driving in Boston. Driving in a crowded city is not the same as driving in a rural area. Driving through the snowy mountains is not the same as driving through a dusty desert or on a winding road along the ocean. If our societal goal is to ensure that autonomous vehicle technology is ultimately capable of performing in all parts of the country, developers must be able to test the technology in multiple states. However, under a patchwork of inconsistent state laws, technology may meet performance requirements in one state and not in another state. Such a situation will impede the ability of a developer to test the same system across multiple states, slowing the development and deployment of the technology. Policymakers should therefore work to promote and advance a single, national framework with appropriate safeguards.

Allowing states to set state-specific vehicle performance requirements for testing could also open the door to state-specific vehicle performance requirements at the time of deployment. Owners or drivers of autonomous vehicles should not be unnecessarily restricted in their ability to travel from state to state, as they can with current vehicles. This is a clear example of why matters of interstate commerce were constitutionally assigned to the federal government. Fifty distinct regulatory frameworks for automated vehicle performance would impede interstate travel and make deployment of a common autonomous vehicle fleet impossible.

Federal Automated Vehicle Policy

By establishing a process that provides consumers and other stakeholders with a level of confidence that autonomous vehicle technology on public roads is safe, the Federal Automated Vehicle Policy ("FAVP") released by NHTSA in September of last year was an important step in cementing federal leadership on automated vehicle policy. While the FAVP was welcomed by the automotive industry, there are several areas of the FAVP that we believe should be addressed before the policy is fully implemented. We look forward to working with Congress and the Administration to address these areas.

First, the FAVP document provides unclear or even conflicting direction to states on their regulatory activity regarding vehicle performance. In Section 2 of the FAVP, which relates to State Model Policy, NHTSA "strongly encourages States to allow [U.S.] DOT alone" [*emphasis added*] to regulate the performance of highly automated vehicles. However, just a few pages later, NHTSA encourages states to require compliance with NHTSA's proposed safety assessment for the technology. While NHTSA has attempted to provide some clarity on this conflicting language, the steady stream of state legislative proposals that regulate vehicle performance indicates that it may not have been entirely successful. We urge NHTSA to clarify in the FAVP itself that it does not intend for states to require compliance with the voluntary safety assessment or to regulate vehicle performance. In fact, a clear and unequivocal statutory or regulatory prohibition on states regulating vehicle performance of autonomous vehicle technology would help to halt or prevent the emergence of a patchwork of state laws.

Second, the FAVP treats testing and deployment of autonomous vehicle technology largely the same. There is an important distinction between developmental testing and public deployment of these systems and the FAVP, particularly as it relates to the applicability of the safety assessment, should appropriately account for that distinction.

For example, under current law, traditional automakers have permission to test prototype vehicles on public roads. However, under the FAVP, all developers - including traditional automakers – are expected to submit a safety assessment for testing autonomous prototype vehicles. We feel that this provision, which seems to conflict with existing law, should be reconsidered.

The submission of a new assessment for each "significant" change or update to the system during testing, as laid out in the FAVP, is likely to hinder the development of the technology. The research, development and testing of these systems is a highly iterative process and involves regular changes and updates. The FAVP should recognize and account for this without sacrificing the transparency that NHTSA seeks. Options for accomplishing this include:

- Establishing a more narrow and targeted safety assessment for testing involving trained safety drivers. An expansive safety assessment could be reserved for testing of systems using members of the public or testing without a trained safety driver in the vehicle.
- Permitting developers to test various features and conduct various phases of testing without the need to submit updated safety assessments with each significant update. This permission could be contingent on the company providing some basic information and assurances that the developer will test these systems responsibly and safely.
- Providing developers the opportunity to submit a general testing plan to NHTSA before commencing testing that could cover various phases of testing for a particular system or feature.

Data Sharing

There has been growing discussion on the need for data sharing in the context of autonomous vehicle testing. There are many goals associated with data sharing. They include sharing data with the government to improve understanding of autonomous vehicle technology, sharing data with the government or public for evaluation of the safety of a particular system, sharing data among developers to help improve the performance of systems, and sharing data for crash reconstruction purposes. While we support the various goals of data sharing, we believe that several important factors should be considered.

If the sharing of data, such as disengagement data, is intended to be used in making a judgment about how “good” an autonomous vehicle system is, it may not accomplish that goal. During research and development, engineers will often intentionally push a system to its limits until it “breaks”. These “edge cases” or “near misses” are specifically sought out to make the systems more robust. Using the number of failures as an indicator of the effectiveness of a system may actually create a perverse incentive to run the technology through easier scenarios in order to make a system appear better to the government and the public. It may discourage developers from testing the scenarios that are needed to make the technology safer.

If the sharing of data is intended to provide “edge case” or “corner case” data that can be utilized by other developers to make their systems smarter, there are a number of important details that need to be worked out. This includes identifying what data should be shared, ensuring that the source of the data is anonymized, deciding where the data will be compiled, and determining who should have access to the data and for what purposes. At the same time, it is important to note that this type of data sharing may not reap the benefits that are intended. Differences between each manufacturers’ sensor configurations may make it difficult to effectively share data, and what might be an “edge case” for one system might be mundane to another.

Therefore, although Toyota agrees with the goals of data sharing, we believe that there is a significant amount of work to be done to ensure that it does not unintentionally delay innovation or worsen safety. We very much look forward to working with other stakeholders to determine how to share data in the most practical and effective manner.

Additional Considerations

Before closing, I would like to provide a couple of additional observations that may prove useful to the Committee on these issues going forward.

First, as previously noted, developing truly reliable autonomous vehicle technology will require extensive testing. The complexities involved in the development, testing, and deployment of autonomous vehicle technology requires a significant amount of public road testing in order to address not only the thousands of traffic scenarios that human drivers will encounter on a regular basis, but also to identify and address as many "edge cases" or "corner cases" as possible. Millions of test-driven miles are necessary, but probably not sufficient, to achieve the reliability that we need for autonomous vehicle technology, particularly if those test-driven miles are through easy or predictable routes. The truth is that all testing miles are not created equal, and developers should be focused on testing scenarios where driving is challenging or even exceedingly difficult. Computer simulation can accelerate and expand the range of testing of these systems, and should – with adequate evidence of validity - be an acceptable equivalent to real-world testing to achieve the billions of test-driven miles that will likely be needed to accomplish this.

Second, it is important that the federal government begin looking beyond testing to deployment of these systems. This includes updating the federal motor vehicle safety standards to address the handful of standards that are inconsistent with or incompatible with autonomous vehicle technology, including those systems that may not require a human driver or human operator in the vehicle. In March of last year, the U.S. DOT Volpe Center released a review of federal motor vehicle safety standards. The report identified safety standards that pose potential barriers and challenges for the certification of autonomous vehicles. While we believe that the report is quite comprehensive, we believe that it would be wise for NHTSA to review this report and carefully consider and solicit feedback on whether there are any other motor vehicle safety standards that could pose a barrier or challenge for autonomous vehicle technology. We also believe that NHTSA should promptly move to update or otherwise address the standards identified in the Volpe Center report, as well as any other standards it identifies as part of its supplemental review.

Thank you for the opportunity to testify before you today. I look forward to your questions.