

SUBMITTED ELECTRONICALLY VIA REGULATIONS.GOV

From: Paul Calhoun - www.rxevolution.me

February 1, 2021

Subject: Framework for Automated Driving System Safety Docket No. NHTSA-2020-0106

Docket Management Facility, M-30
U.S. Department of Transportation
West Building, Ground Floor, Room W12-140
1200 New Jersey Avenue SE Washington, DC 20590.

To Whom It May Concern:

It is with great anticipation that I see this notice of proposed rulemaking. The need for meaningful AV regulation has grown from matter for individual States to a significant national concern. After reading the ANPRM, I believe that responses to the questions asked at the end will encompass most of my comments, so I will begin with those. I'm doing so out of order because I believe that my vision and concerns start from a different place than most, that of a robotics expert looking at regulation rather than a regulator looking at an autonomous system.

Questions:

Question 6. Do you agree or disagree with the core elements (i.e., "sensing," "perception," "planning" and "control") described in this document? Please explain why.

Question 7. Can you suggest any other core element(s) that NHTSA should consider in developing a safety framework for ADS? Please provide the basis of your suggestion.

I believe that the core elements conflate two major elements together, namely "Planning" and "Prediction." To take a hypothetical pipeline: Sensing perceives an object and sends data about it to Perception. Perception classifies the object. Prediction determines where the object will

be based on its classification and both past and ongoing data from Sensing. Planning determines how to modify the existing path plan and behaviors based on that prediction and classification. Control executes the new plan.

I would propose that Prediction be raised to the same level of importance as the core elements laid out in the ANPRM. Looking at the fatality of a pedestrian in Tempe, the failure was as much one of Prediction as Perception, but if you take Prediction out of it, not of Planning, Sensing, or Control. The pedestrian's classification changed repeatedly, and the Prediction subsystem failed to make use of existing data to predict the pedestrian's path. The Prediction subsystem, in fact, was configured to only use prior data for as long as the Perception subsystem maintained a single classification, so it was effectively throwing out valid data because Perception had failed. Had the Perception subsystem settled on a single classification *or* the Prediction subsystem used data based only on past motion, it is likely that the vehicle would have stopped much sooner. A good Prediction subsystem does not require Perception to make a final decision on the nature of an object to determine its likely path; it can begin that process using the object's velocity as observed.

A preliminary classification of *solid object* would have also been a significant improvement, though this does not impact the core elements.

Prediction isn't just determining where objects have been and will be. It's determining where they have been and will be *relative to the vehicle in motion*. That's predicting both the motion of the object and the motion of the vehicle with respect to them, which can be a separate calculation that goes beyond using the Planning to determine future states of the vehicle.

I would add an element that is not considered here at all, that of Communication. Communication, whether it be to other vehicles, humans around the vehicle, or infrastructure, will be a crucial part of AVs in the future. While it is not part of the pipeline used for vehicle behavior in a vacuum or on a test course, Communication is key for both the vehicle's overall environmental awareness and for the awareness of stakeholders regarding the vehicle's intent. To leave Communication out of the equation is to diminish rather than improve overall vehicle capabilities, since humans do everything from listen to the radio for road closures to

gesture at pedestrians to clarify their intent at an intersection. Even if the form of this Communication is different, the overall result must remain at least as good.

Questions:

Question 1. Describe your conception of a Federal safety framework for ADS that encompasses the process and engineering measures described in this document and explain your rationale for its design.

Question 2. In consideration of optimum use of NHTSA's resources, on which aspects of a manufacturer's comprehensive demonstration of the safety of its ADS should the Agency place a priority and focus its monitoring and safety oversight efforts and why?

The prior answers to questions 6 and 7 reveal a way of looking at AVs which differs significantly from how automobiles have been considered until now. AVs are vehicles, but they are also *agents*. It is not enough to put them through the same type of testing as existing vehicles because those types of tests would, at best, test the Sensing and Control elements with a minimal meaningful testing occurring on Perception, Prediction, and Planning. We must think of AVs as both vehicle and agent, car and driver.

Let's start with the driving exams and tests given to those who are to be licensed. Something analogous must be administered to an AV. The benefit is that a more comprehensive test can be given to an AV than a human because the answers given and behaviors shown by one AV applies to all AVs of the same configuration and software version. Realistically differences in sensor performance and minute differences caused by physics will mean that testing a sample of a 5-10 AVs in a single configuration will be required.

There are several exams and tests given depending on the individual State. I will take one of the more stringent set since they will show the way for something as important as testing an entire fleet. First, testing the comprehension of the rules of the road, road signs, etc. Second, the ability to *perceive* when those rules are to be applied (ex. eye test, hearing test). Finally, the practical test to see how the driver is able to cope behind the

wheel, first in a controlled testing course, and then on a representative section of roadway.

How does give a driving test to an AV? We need to break down these tests and understand an important underlying assumption. The reason why a combination of written test and test of the senses (again ex. eye test) is that humans fall into a certain cognitive space. If we can see a road sign and we know what it means, we are assumed to be able to react in the correct manner. The practical test checks this with a basic road sign, usually a stop sign. This works because the majority of human brains fall into an understandable and bound cognitive space. We think alike because our meat brains are the same layout.

An AV does not have the same cognitive space as a human, so our tests don't work on them as well. A Perception element is not a visual cortex; it can't generalize nearly as well *but* it does function much more predictably. The cognitive space of an AV is bounded by the hardware and mathematics which underly its algorithms. Thankfully, the vast majority of hardware in an AV evolved from Von Neumann architecture, and the vast majority of algorithms evolved from backpropagation, Bayesian probability, etc. It is not a perfect match, but it gives a helpful basis in understanding where to begin with testing an AV *by itself*.

The *by itself* is important, because the AV as a single unit is not the only matter being considered, but it's the first thing we can look at because as a self-contained unit, we can think in terms of being performance based and much less about prescriptive regulation.

In terms of the (singular) core elements of Sensing, Perception, Prediction, Planning, and Control, the closer to the middle of the list, the more important the test becomes. That is, the core elements in order of importance would be something like Perception, Prediction, Planning, Sensing, Control. This much can be understood from accidents and fatalities so far observed, and simply because testing of sensors and control are already mature. Any vehicle currently on the road has a well-built Control system, and testing of Sensor performance is easily done as a unit test. *Sensor Fusion* perhaps less so, though that is under the umbrella of Perception, which is highest on my notional list of importance.

If we think of it from the perspective of a student driver going in for their license, we need to make sure they can see the pedestrian, know what to do when one is in the crosswalk, and guess with acceptable accuracy where they will be *relative to the vehicle* at any given interval. Or perceiving a road sign, knowing how to obey it, and then determining how obeying that sign will affect future states of the vehicle.

My highest and equal priority would be on the following things:

1. Accurately determining a range of objects in the vehicle's view
2. Accurately determining those objects' motion - if any.
3. Accurately determining what the vehicle should do in relation to the object to maintain safety for all stakeholders.

There may or may not be a set range on any of these. That is, a set time in which the vehicle should reach X confidence rating (confidence ratings being entirely subjective to each algorithm is a significant issue which requires a standard), a set distance from said object, or a set time/distance by which the vehicles replans safely. This is because of the Control aspect. Different vehicles have different stopping distances, and that should be accounted for in determining if the vehicle makes its decisions in a safe time frame. The most important result is that the vehicle processes all data and decides in time for Control to safely execute.

However, a good result based on bad information should not be a successful test. It must take the correct action for the correct reason, or else it is highly likely to take incorrect actions later. One of the benefits of AVs over humans is that we can look much more critically and accurately at *why* the system behaved the way it did.

With all that set up, it's time to turn outwards to the aforementioned Communication element. It's not enough that the vehicles can take the correct action in any given scenario, because many scenarios require communication with an external stakeholder. Whether it's retrieving data after a crash, emergency highway lane redirections (ex. all lanes become outbound in a hurricane), or communicating the vehicles' plans to a pedestrian in a curb cut. All these things require robust V2X

communication in various forms. Being able to send and receive this data will be an important part of testing going forward, and will likely be the first place where specific regulations will be required. Performance-based testing is fine when the data involved exists only within the AV in question; how they move information around within their system is their own decision. Once it needs to leave the boundary of the system of interest, a standard is needed.

Much like Open Mission Systems (OMS)¹ in the Department of Defense, we need an Open Driving Systems (ODS) standard for V2X communication. This isn't a major revelation; it's been the subject of papers for years. It will be DoT's role to determine and require the *minimum* information that passes in these messages. It's not the purview of this comment to give an exhaustive list, but examples would be routine and emergency road modifications like closures, position data shared with other AVs, and the data exported for crash investigations.

There is also the *legibility* and *predictability* of the AVs, which is another branch of Communication, this time between the AV and external humans. *Legibility* is understanding the intention of the AV (ex. turn indicator lit at an intersection). *Predictability* is knowing how it will execute that intent (ex. knowing that the AV will remain in the same turn lane, or cross turn lanes partway through the turn because it will need to make another turn soon after). Regulating this is an extension of existing standards and regulations on elements like the color, position, intensity, and frequency of turn indicators, and similar requirements on brake lights. While some of this will merely be continuing to use existing communication methods (ex. turn indicators and brake lights), humans often look at the eyes, expression, and gestures of drivers to determine their intent. An analogous method must be incorporated in AVs to maintain the same level of legibility and predictability. Tests have been done using displays made to appear like eyes on the front of vehicles so that pedestrians know that the vehicle is aware of their presence².

¹ <https://www.vdl.afrl.af.mil/programs/oam/oms.php>

² <https://www.dezeen.com/2018/09/04/jaguar-land-rovers-prototype-driverless-car-makes-eye-contact-pedestrians-transport/>

Questions:

Question 4. How would your framework assist NHTSA in engaging with ADS development in a manner that helps address safety, but without unnecessarily hampering innovation?

Question 5. How could the Agency best assess whether each manufacturer had adequately demonstrated the extent of its ADS' ability to meet each prioritized element of safety?

By keeping singular AVs to performance-based testing modeled on driving license exams, innovation can be focused on results rather than satisfying specific numbers that will likely change based on improvements in technology. Mandating and testing a unified V2X communication framework has minimal impact on innovation since it requires at most a set of middleware components to communicate with the proposed ODS, and at best *helps* AV developers focus on important safety-related data.

The question to ask is *can this AV pass a driving test?* The key part of this is to remember the difference in cognitive space, and the similarities. Testing AVs using the exact same conditions could result in AV developers using simulations of those courses as training data; and a maxim in machine learning is never to test with training data. Similarly, a driving instructor doesn't test their student using the actual driving test. It produces local solution spaces in computers, and complacency in human drivers. Regularly changing the parameters of the course helps with this, and having multiple courses. Test one vehicle in each AV fleet on all the courses, and change them regularly. If they pass them all, that gives a high confidence that they will be effective in normal driving conditions at least.

Unusual driving conditions can be tested, but will require specialized courses. State involvement may be helpful in this. As each State has its own DMV and driving test, so too can they test AVs for their individual State, increasing the variety of conditions tested across the fleet. This will include not just the testing of the singular AV, but of its communication as well. At a minimum, its communication with external humans will be tested as part of any normal driving test, but also its ability to ingest and comprehend the specific laws of the State and municipality, which must either be communicated using V2X or by looking up those rules based on the vehicle's location in an internal database.

Questions:

Question 14. What additional research would best support the creation of a safety framework? In what sequence should the additional research be conducted and why? What tools are necessary to perform such research?

I recommend focusing research on creating a meaningful ODS, and on developing a version of the driving test which accounts for the cognitive space of AVs. Ideally running concurrently, but realistically it's more important to get the test working than the communication standard, which may emerge from industry over time anyway, especially as municipalities levy their own requirements about how much data the vehicles must provide and accept.

The tools required are best determined by a partnership between government, industry, academics, and the public. There are well prepared Human-Robot Interaction (HRI) departments at universities like Carnegie Mellon's Robotics Institute and at Georgia Tech's Robotics Lab. They would be well equipped to investigate how to translate a human driving test into the machine domain. Community engagement will also assist in this, as local people can raise issues that would otherwise get lost in the race to a broad-brush solution.

Questions:

Question 16. Of the administrative mechanisms described in this document, which single mechanism or combination of mechanisms would best enable the Agency to carry out its safety mission, and why?

Question 17. Which mechanisms could be implemented in the near term or are the easiest and quickest to implement, and why?

For individual AVs, starting with a mix of voluntary mechanisms and regulatory mechanisms would be best. Keep the AV makers in the loop as the driving tests are developed, working with them to determine the best way forward in what they should include. Elements like the list of objects which must be classified correctly (ex. pedestrian, bicyclist), which will change as AVs become more capable and evolve with the aid of the developers. Having that guidance and framework will keep the

developers moving towards a safety-conscious goal as well as contributing to that same goal.

The developers can say best what *can* be tested, and regulators what *should* be included to assure safety. Finding common ground and collaborative development of the test structure will make sure that the AVs are safe as the technology becomes more mature and the mechanisms more regulatory. Developed capabilities will build the runway for effective regulation.

Required reporting is crucial, even if the specifics of the data must remain confidential. The overall results of analysis of the data should be made public and it's in the interests of the developers to have a common agency collecting the data and publishing the results of analysis. Being able to attain a safety ranking from DoT will spur safety innovation and increase public confidence in the reliability of the AVs that score well.

AV developer collaboration on the ODS will also be critical, and the implementation of a draft framework as regulatory requirement will also be very important. In the same way as the testing of individual AVs will build regulatory runway as technology develops, so too will the development of the standard build runway for required use of the standard. Something like it being voluntary for the first two years of release, and then each release makes the earliest voluntary release into a mandatory release. That will give them time to become compliant as well as the opportunity to get ahead and shape the direction the standard takes.

Another place where research and stricter regulation will be required is in emergency disengagement. AVs in the L2 and L3 range have a *control inequality* where the AV can disengage at any time, making the driver responsible for what occurs and how the vehicle behaves, but the driver cannot go the other way. This is, of course, because the driver *is* responsible and theoretically more capable. However, it can also be an easy way for the AV and the AV maker to avoid liability and responsibility for adverse effects. The DoT regulations must require a minimum disengagement window in which a reasonably effective driver can undertake meaningful emergency maneuvers. The AV must "take ownership" of its situation within its ODD, and either be able to handle

emergency behaviors within that ODD or recognize the need for them within a time frame which allows for the human driver to take control and make a meaningful effort to avoid a collision or other adverse situation. An AV disengaging within the window should be treated as if the AV didn't disengage at all, and full responsibility be placed on the AV and its maker.

One of the thorniest issues is the current software development cycle. Many companies practice sloppy DevOps and few practice DevSecOps at all. The pipeline is based on a form of Agile which promotes releasing new versions over comprehensive testing and documentation. This has been an issue in the past and will remain one in the future. Model Based System Engineering (MBSE) has a major automotive component in its development, and most of the large legacy automakers continue to use it extensively. Newer automakers and AV developers do not practice MBSE or likely any significant documentation practice at all. While testing in simulation is a big component, *developing* in a model may be less so. A cursory inspection of job offerings shows MBSE in constant demand from major Detroit automakers, but almost totally unseen among the AV developers.

Discussion of this aspect swiftly moves into the murky waters of monopoly. Is it detrimental to innovation and good design for the same developer to produce Perception, Prediction, Planning, and Control? It would improve both vehicle communication standards as well as interoperability, documentation, and robustness of subcomponents if different developers each specialized in one of those elements. Mandating this in any meaningful way is well outside my understanding of the purview of DoT. However, it is reasonable to require more robust software release strategies in the form of requiring unit test, regression test, and safety critical design. The choice could be put before the developers: either demonstrate good DevSecOps and a robust digital thread, or every release no matter how small will require a *separate* and *full* test of the vehicle fleet.



Beyond the issues raised by the Questions in this ANPRM, there is the matter of the multiple layers of AVs currently being road tested. My response to the questions mostly address private motor vehicles which will be on public highways, since they will make up the bulk of what will require testing and safety analysis in the near term.

Near-term and mentioned in the ANPRM are low-velocity shuttles. I envision a minimal testing requirement for them since they will be driving a set route with minimal variation from day-to-day, and thus be lower “volatility” in terms of safety issues and need for changes. If they can drive the route they are set to one day, they will likely drive it just as well from then on. Testing them *in situ* with a few standard elements such as a simulated pedestrian, bicyclist, and other vehicles will likely suffice for long-term certification of their readiness.

Near-term, *not* mentioned in the ANPRM, are a big problem. Low velocity delivery crawlers. Tracked, wheeled, and sometimes legged robots that carry one or more parcels the “last mile” between depots or small businesses and customers. Most will not be on roads, which makes them an even bigger problem. They’ll be on *sidewalks* and *in crosswalks*. How they communicate with pedestrians and navigate these often-clogged byways without becoming a danger or a nuisance may or may not be a DoT issue, but a framework for States and municipalities to follow would be a helpful effort. Already there has been an incident of one of these crawlers putting a wheelchair bound pedestrian (vulnerable stakeholder) in danger by blocking her access to a curb cut while she was in a crosswalk on a busy street³. This may seem like a local issue until one considers that the Federal government does have jurisdiction over some sidewalks on their own land, which means there will be a Federal agency that will need this guidance to regulate Federal sidewalks and land. If nothing else, this is an ADA issue.

Finally, long-term, there’s the concern about trucking. Not just the near-term carrying of the usual cargo by 18-wheelers, but autonomous trucks pulling hazardous materials. Granting an algorithm a Class A CDL

³ <https://nationalcenterformobilitymanagement.org/what-can-we-learn-from-emily-ackermans-fight-with-a-sidewalk-robot/>

is a decision that will take a great deal of consideration, and we would do well to start considering it right now.

.....

I would like to thank you for your attention and for taking steps to assure us of a safe American rollout of AVs as they appear in our cities and towns. No doubt many comments were sent in by organizations attempting to gain your attention to sell their product. In that way, my comment is no different. I am aware of the brain drain in the Federal workforce, as well as the difficulties being faced as more civil servants retire than are hired each year. Since USAJOBS seems to have some difficulty with hiring, I thought I'd go directly to the source with this comment. I have a deep interest in public policy and autonomous systems in general. I have a Master's of Science in Robotic System Development from Carnegie Mellon, and would be very interested in discussing a position within the government if DoT finds itself in need of more personnel in this area.

V/R

Paul Calhoun

<https://rxevolution.me/contact/>