Automated Vehicles: How Safe is Safe Enough?
October 16, 2018

The connected and automated vehicle industry is growing rapidly. Last year, a study from Intel and Strategy Analytics estimated that driverless vehicles could represent $7 trillion of economic activity annually by 2050.[1] Achieving this market size will require each manufacturer to answer critical questions about whether its product is safe enough. While a consensus answer has yet to emerge, our history of development in transportation suggests that it will be a function of three key criteria: regulatory demand, consumer acceptance, and a manufacturer's tolerance for risk. Above all, stakeholders must have confidence that every automated vehicle on the road reliably performs to the agreed upon safety performance standards or negative societal, individual, and business consequences will be the result.

Traditionally, an automotive manufacturer creates a testing program for a new vehicle by looking to regulations and their experience—drawing upon a variety of scripted scenarios that have historically been useful predictors of vehicle safety and coupling these with reliability and performance testing to assess the overall safety of the product. With an automated vehicle, these scripted scenarios do not exist, and the idea that a vehicle design has a set of performance and safety characteristics that can be assessed once while the vehicle is "new" and thereafter remain constant throughout its design life is flawed. The recent surge in the development of automated vehicle technology is taking us toward automated driving systems (ADSs) that are not only designed by exposing the algorithms that control system behavior to curated collections of training data and having the algorithms learn from these data—but to ADSs that continue to learn throughout the entire product lifecycle and across entire deployed fleets. This leads to the open question of whether and how frequently to re-test the automated vehicle as its ADS is incrementally modified: At what point do the original test results no longer apply?

It is tempting to think that a vehicle under the control of a continuously learning ADS that successfully navigates a certain number of miles travelled is on the right path to demonstrating a "safe enough" level of performance. But this metric, while easy to calculate, can be a weak predictor of future performance. The ADS in an automated car is entrusted with making the correct decisions to navigate a vehicle from point A to point B; however, it is the quality, quantity, and type of decisions made in this process that reflect ADS performance. Instead of a digital metric indicating success or failure while navigating the user-selected (or algorithm-selected?) route, a more appropriate approach may be to borrow an idea from risk analysis and look at the exposure a vehicle experiences while travelling from A to B, as well as the margin available to the system while navigating that risk. Instead of vehicle miles travelled, consider the following:

- What kind of interactions/decisions were experienced/ made? Were there any "edge" cases experienced, or did the ADS only experience "conventional" scenarios?

- How many interactions/decisions were experienced/made?

- With what margin or opportunity for future action was each decision made?

---

[1] https://newsroom.intel.com/news-releases/intel-predicts-autonomous-driving-will-spur-new-passenger-economy-worth-7-trillion/

By looking at each of these three questions, it is possible to gain insight into the coverage of the operational domains that the testing/training covers as well as the robustness of the performance estimate. When it comes to safety, consumers are likely to trust an automated vehicle when it performs significantly better than their own driving abilities. It is not enough for a vehicle to consistently turn right or left and stop at traffic signals when needed. Can the vehicle turn left onto a street that is hosting a parade and adjust accordingly? Can it adapt to a malfunctioning stoplight or follow a police officer's instructions that are contrary to what the light is indicating? Manufacturers of automated vehicles must demonstrate that their products can mitigate safety risk across a myriad of real-world interactions. Failure to do so can jeopardize consumer trust, limit commercial viability, and impact human lives.

Whereas other industries—or even the conventional transportation industry—may look to well-defined regulatory standards for guidance, the automated vehicle industry is immersed in a rapidly evolving, piecemeal regulatory framework. Automated vehicle regulations currently vary from state to state with limited consistency across locations and only a very general federal strategy for harmonization. As a result, innovators must take the first step in self-certifying that their products are safe enough to be on the road. Manufacturers must be able to say, "I know where the bar is regarding safety, and here is where I am with respect to that bar." This requires a strong understanding of the different regulatory environments within the United States and prospective global markets.

Manufacturers of automated vehicles often partner with third parties to help them navigate the above complexities. Exponent offers a fully integrated team of engineers, human factors specialists, and regulatory experts that can help manufacturers answer the question of "How safe is safe enough?" With over fifty years of safety experience, start-ups and traditional vehicle manufacturers alike can find a partner in Exponent for their automated vehicle needs.

John D. Pye, Ph.D., P.E.
Group Vice President, Office Director & Principal Engineer
Phoenix
(623) 587-4143 | jpye@exponent.com

E^xponent®