The Brave New World of Driverless Cars

M.L Cummings
Director, Duke Robotics

Recently the National Highway Traffic Safety Administration (NHTSA) released a set of federal guidelines that outlines how manufacturers should test and deploy self-driving cars (*1*). Shortly thereafter, the California Department of Motor Vehicles (DMV) released a revised draft of regulations that essentially opened to door for manufacturers to test cars with no human in the car on public roads as long as manufacturers abide by the federal government's 15-step assessment guidelines[1].

Taken together, particularly in light of similar announcements by many cities including Boston, Pittsburgh, and Austin, it is clear that driverless cars are a near-term reality. Soon, the general public will be driving in the same environments as driverless cars engaged in testing. Once these cars theoretically generate enough miles to convince both state DMVs and NHTSA that they are safe enough, then everyday consumers will either be able to buy their own driverless car, or summon a rideshare driverless car.

If these cars work as advertised, such a goal is laudable as vision and mobility impaired people will have new forms of transportation that could be life changing, and the annual ~38,000 fatalities caused by driver error could be significantly reduced. But such claims of benefit should also be interpreted in light of the significant amount of money that is being poured into this industry. Global automotive research and development expenditures in 2016 are estimated to be 94.2 billion dollars, which is 3 times that of 2016 global aerospace and defense expenditures of 30.4 billion dollars (*2*).

Industry generally paints a very optimistic picture of when driverless cars will be ready for widespread deployment, anywhere from 2018-2025[2]. However, two recent fatalities of Tesla drivers in China and the United States, while driving using a feature known as "Autopilot", have highlighted the complexities and frailties of these semi-autonomous systems. These fatalities, as well as other related accidents, call into question both the capabilities of the systems, as well as the lack of a testing and certification process.

In the first Tesla fatality in China, a driver slammed into the back of street cleaner on a highway while in autopilot mode. Five months later, a driver was killed in Florida when his Tesla, also in autopilot mode, failed to detect a tractor trailer turning in front of it, and hit it broadside at 74 mph[3]. Tesla insists that both drivers were at fault for not paying attention to the autopilot, which is technically a "driver-assist" technology and not intended to be a fully autonomous car.

---

[1] https://www.dmv.ca.gov/portal/wcm/connect/211897ae-c58a-4f28-a2b7-03cbe213e51d/avexpressterms_93016.pdf?MOD=AJPERES
[2] http://www.driverless-future.com/?page_id=384
[3] http://www.nytimes.com/2016/09/15/business/fatal-tesla-crash-in-china-involved-autopilot-government-tv-says.html

Unfortunately, both drivers did not understand this nuance, which is no doubt exacerbated by calling the technology an “autopilot.” Moreover, such examples of “mode confusion” are well known in aviation and will only increase as automation becomes more prevalent.

These two fatalities highlight several issues I raised before the Senate Commerce Committee in March of 2016 (*3*), namely that while there are many unknowns that we will discover as driverless cars enter the market, there are many known factors that are not being addressed by manufacturers. For example, Tesla knew about the inability of the autopilot to detect static objects on highways, and it even warned drivers in its owner’s manual that the car may not brake for stationary vehicles, especially while driving over 50 mph[4]. A combination of a significant flaw in the perception system of the car (how it ‘sees’ the world) and the lack of transparency to the drivers led to these fatalities, and these problems are not easily solved.

The sensors on driverless cars that help them see can include some combination of radar, LIDAR (Light Detection and Ranging), computer vision, and ultrasound devices. No single technology can provide complete coverage so some combination of these sensors must be used, which requires complex data fusion. Moreover, each of these sensors has known limitations, as illustrated by the Tesla fatalities, and inclement weather, including fog, rain, and snow presents additional problems. Moreover, post-processing of the data gathered by these sensors requires significant estimation and pattern matching, often referred to as machine learning, so when an expected driving scene does not match what the observed scene is by the sensors (which may themselves be flawed), an autonomous car may not be able to reason accurately about the world around it and what the correct next actions are in the required time.

These sensor and post-processing difficulties are widely known in automotive robotics and indeed, across all robotics industries that rely on such technologies (including unmanned aerial vehicles, aka drones, manufacturing robots, and medicine). While significant effort is underway in both academic and industry research environments to improve these technologies and processes, there is substantially less effort in determining how to develop testing strategies to ensure these stochastic systems can work both in expected driving conditions as well as the boundary conditions where catastrophic failures happen.

Because of the strong reliance on pattern recognition and probabilistic reasoning in driverless cars, previous test strategies used for deterministic systems simply do not work. Because of the embedded complexities of such stochastic systems, these cars, for example, will not compute a solution to a four way intersection the same way each time it occurs. There is no industry-wide

[4] http://money.cnn.com/2016/07/07/technology/tesla-autopilot-name

consensus, either for driverless cars or for any unmanned systems, on how such probabilistic systems can and should be tested to guarantee any level of safety.

Currently, driverless car companies have been generating a “miles driven” metric to provide assurances of safety. RAND has stated that driverless cars must drive 275 million miles without a fatality to ensure these cars are as safe as human drivers, with 95% confidence (*4*). Tesla only logged 130 million miles before the US fatality (the most of any company), so Tesla fell significantly short by this metric.

Miles driven is simply not an acceptable solution for demonstrating that a technology is safe for public roads, especially when such numbers were generated in sunny climates with white lines clearly visible on well-maintained highways. In addition to the need to test the stochastic reasoning of these autonomous systems, the cars must be tested at the corner cases, which are the worst possible scenarios that could be faced by these cars including snow, ice, fog, dense pedestrian and bicyclist environments, unexpected behaviors from other cars, etc.

In its recently released federal guidelines, NHTSA lays out a 15 point assessment plan that states should follow in order to assess whether driverless car technology is ready for use on public roads (*1*). This assessment plan only addresses very high level areas of concern like privacy, system safety, and object and event detection and response. It does not provide any specific guidance or assistance for how to assess each of these 15 areas, leaving each state to interpret and perform their own safety assessments, which no doubt will vary widely.

The evaluation of driverless cars is extremely difficult and requires engineers who are both hardware and software experts, especially those in artificial intelligence. In its federal guidelines, NHTSA admits that it does not have on staff those people qualified to make such assessments, and suggests that it *may* develop a network of experts to help it in better understanding these issues.

Unfortunately, no other plans are made to centralize or disseminate any expert knowledge, at either the federal or state level. This means that in a very short period of time, state governments will be expected to acquire the expertise to assess driverless car test plan validity and comprehensiveness, given that there are no commonly accepted standards or even any consensus on how such testing should be done. Given that the area of test and evaluation of autonomous systems is nascent with very little research, either theoretical or empirical, it s a tall order to now expect state governments to do what researchers have not yet demonstrated.

Given the lack of a principled approach to testing of autonomous systems, it is not clear what the implications are to the general public. As mentioned previously, California will soon let driverless cars be tested on public roads with no driver in the car. However, a remote operator will be monitoring the system. This raises an important issue of informed consent for the public. Given that such tests will be state and federally sanctioned through NHTSA’s 15 point plan, an important issue not addressed by NHTSA in their guidelines is whether 45 CFR 46 will apply, which mandates that all humans involved in an experiment should explicitly give their consent to be involved.

The basic question is whether drivers should be given the option to share the road with one or more driverless vehicles undergoing testing, especially with no safety monitors physically in the car? These cars have no established minimum safety standards, and their determination of road worthiness and public safety are left to the discretion of state evaluators who likely do not have the appropriate background to make such a judgment. At a minimum, there should be a discussion of how to clearly mark such cars so that drivers sharing the road with driverless cars undergoing testing have some understanding of the test environment they did not volunteer to be in.

So what do all these issues mean to the research community going forward? The promise and potential benefits of driverless cars could be transformative, but further research and development is badly needed as the rush to deploy driverless car technology has outpaced our technical underpinnings. As a result, we need significantly more research in a wide number of areas including sensor development, improved artificial intelligence and machine learning approaches, test and evaluation of autonomous systems, and research into the legal, ethical and public policy implications of driverless cars.

We also need significantly more interdisciplinary work across these fields in order to communicate the capabilities and limitations of these probabilistic systems. For example, we need much more work in an area known as explainable AI (artificial intelligence), so we understand how to better communicate the outcomes of machine learning algorithms to both researchers and policy makers. Moreover, the recent Tesla fatalities highlight the gap between engineers who design these complex systems and humans who do not understand them.

Much more research is needed in human-robot interaction such that reciprocal intent is effectively communicated between all entities in driverless car sociotechnical systems including the cars themselves, human operators, pedestrians, bicyclists, etc. Broader sociotechnical questions should also be addressed like what the impact could be on public transportation, as well as what fuel types and requirements are being considered and how projected demand could affect air quality and overall congestion.

Universities and colleges also need to increase the numbers of students going into these fields. The demand for electrical, mechanical, and computer engineering students and software developers, all core to the driverless car community, exceed supply (*5*). Because of the growing sociotechnical issues, programs that address the interdisciplinary aspects of driverless cars should be developed, as those graduates are badly needed by government and industry. Universities need to adapt to these growing demands, as well as government agencies and foundations that provide scholarships and incentives for relevant new programs.

Such change and growth in educating a multidisciplinary robotics workforce is critical both in the US and worldwide as driverless cars represent just one rapidly growing robotics industry. Commercial drones, manufacturing robotics, medicine, and other industries attempting to insert more autonomy into their operations are competing for the same people, and the chokepoint currently resides in higher education.

1. *Federal Automated Vehicles Policy:Accelerating the Next Revolution In Roadway Safety* (2016).
2. Industrial Research Institute, "2016 Global R&D Funding Forecast," (R&D Magazine,, 2016).
3. M. L. Cummings, in *Senate Commerce Committee*. (US Senate, Washington DC, 2016).
4. N. Kalra, S. M. Paddock., "Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?," (RAND, Santa Monica, CA, 2016).
5. Y. Xue, R. Larson, in *Monthly Labor Review*. (Bureau of Labor Statistics, Washington DC, 2015).