AUTONOMOUS VEHICLES STILL DECADES AWAY: 2019

Raaga Kannan
Thayer School of Engineering
Dartmouth College
NH, USA

Ronald C. Lasky, Ph.D., P.E.
Indium Corporation, Dartmouth College
NY, USA
rlasky@indium.com

ABSTRACT
The automotive industry, once predicting the arrival of fully autonomous vehicles by 2020, has backed down from such optimism as industry experts recognize the difficulties of bringing level 5 automation into the hands of consumers. This paper will first introduce the stages of automation defined by the Society of Automotive Engineers (SAE International). Secondly, it will examine the challenges required to progress from existing advanced driver-assistance systems (ADAS) to level 4 and 5 autonomous vehicles. This section will focus on the need for higher precision sensors and software standards as well as the development of cognitive functions such as perception in existing software to navigate daily traffic patterns encountered by human drivers. The inability of current AI technologies to accomplish such a feat will then be discussed.

Next, assuming automakers successfully develop the necessary technologies for autonomous vehicles, difficulties of testing the safety of such vehicles will be addressed.

This paper will conclude with a discussion of the dangers of releasing level 3 autopilot systems to consumers.

Key words: autonomous driving, artificial intelligence, autonomous vehicle testing and safety

INTRODUCTION

While Elon Musk remains confident of his promise to deliver to customers fully autonomous vehicles by the end of 2020, industry leaders are admitting that a driverless future is further out than originally anticipated.

A recent Design News article quotes John Krafcik, CEO of Google's self-driving car unit, Waymo, stating that "It's really, really hard… Autonomy will always have some constraints."[1]

Jeffrey Funk, former professor at the National University of Singapore, describes the new wave of technological hype as a resurgence of "irrational exuberance", a term coined by former Chairman of the Federal Reserve, Alan Greenspan, to describe the dotcom bubble in 1996.[2] As journalists and media companies are incentivized to produce more favorable forecasts to attract viewers, the result is news like BBC's 2018 article titled "Why you have (probably) already bought your last car." [3]

Now, automakers and suppliers alike are acknowledging the difficulties of designing vehicles that can assess surroundings and accurately make decisions in non-controlled environments. While companies, including Ford and Waymo, have arrived at this conclusion drawn by engineers years ago, it is a wonder whether Musk is just strapped for cash or already asleep at the wheel.

Although fully autonomous vehicles may not be on the road for a few more decades, this brings good news for manufacturers of electronics, as there will be many electronic advancements required before fleets of robo-taxis can be seen on the roads.

STAGES OF AUTOMATION

Over the years, vehicle safety features have evolved from seatbelts and classic cruise control systems to advanced safety and driver assistance features including lane-keeping, emergency braking, and adaptive cruise control.

One source of false expectations surrounding the arrival of autonomous vehicles is the ambiguity of terms such as "self-driving" and "autonomous" when used in media. SAE International clearly defines the degrees of automation using six levels, with 0 defined as no automation, and 6 being full automation as shown in Figure 1.[4]

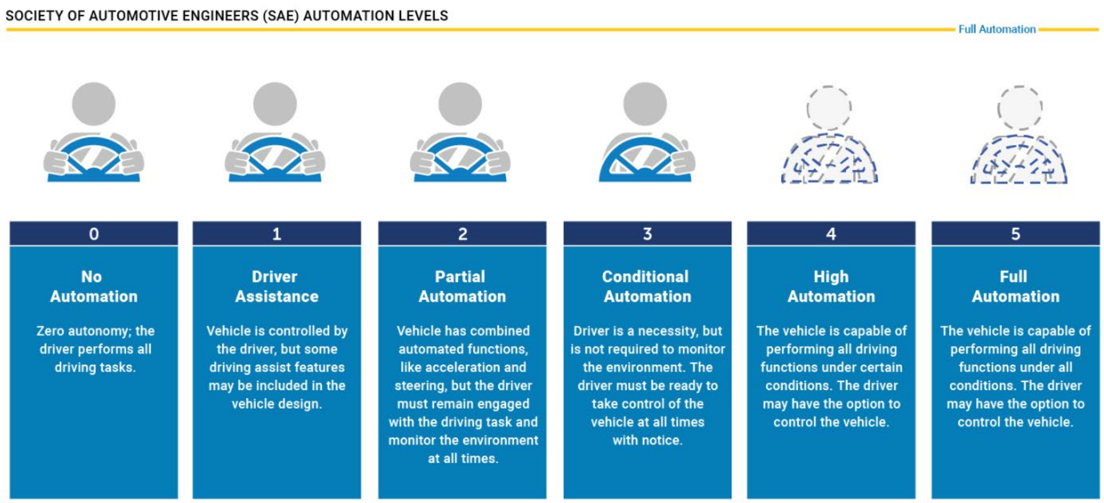

Figure 1. Society of Automotive Engineers (SAE) Automation Levels. This chart is from the National Highway Traffic Safety Administration.

According to SAE International designations, level 0 is no automation where the individual must complete all driving tasks. With level 1 automation, ADAS systems assist drivers in steering or accelerating/decelerating, but not both. A level 1 system would be a vehicle with either adaptive cruise control or lane-keeping, for example. With level 2 systems, the vehicle can perform both steering and accelerating/decelerating; however, constant human monitoring and action in all other tasks is required. A vehicle with both lane-keeping and advanced cruise control is an example of level 2 automation such as Tesla's autopilot feature.

With level 3 automation, the vehicle now performs all tasks of driving in specific circumstances, but it demands that the driver maintain attentiveness and take control in the event of an error. Despite the misleading name, Tesla's latest full self-driving capability (FSD) feature is an example of a level 3 system, providing automation on freeways under optimal conditions with the requirement that drivers keep their hands on the wheel. Tesla's "feature-complete" FSD, announced for potential release by the end of 2019, allows for automated driving through city streets as well. Even with the additional capabilities, this system still classifies as level 3 as it requires that drivers are prepared to take control at any time.

Level 4 and 5 automation are increasingly complex as the human driver is no longer required to monitor or serve as a back-up. With level 4 automation, the vehicle performs all driving tasks in restricted conditions. Level 5 automation allows the individual to relinquish control in all situations and act as a passenger. Level 5, or full automation, is what is referenced by Musk when he announces that customers will be safe to sleep in their vehicles and wake up at their destination by 2021. Unfortunately, the path from ADAS to autonomous driving (AD) is not linear and requires overcoming obstacles involving the performance and testing of such vehicles.

REQUIREMENTS FOR AUTONOMY

While some look at the over 50,000 autopilot enabled Teslas as evidence for the imminent arrival of the level 4 and 5 autonomous vehicles in 2020, the highly advanced level 2 Autopilot system does not provide an iterative path to the development of fully autonomous vehicles. Carlos Holguin, co-founder of AutoKab, recently quoted in an EETimes article, "As the famous saying goes, 'the electric light did not come from the continuous improvement of candles.'"[5]

Transitioning from ADAS to AD requires even greater precision in the sensory devices used (radar, LiDAR, GPS, and cameras) as well as the software standards itself as errors in detection and action are no longer the responsibility of human drivers to correct.

What primarily differentiates ADAS from AD, however, is the new software requirements for vehicles to handle daily road encounters that drivers face. Sensory capabilities alone are no longer sufficient. Using sensors to determine speeds and distances of objects is much easier to accomplish, which is why systems such as cruise control have been around since 1958.

AD requires that vehicles not only sense their surroundings, but also perceive the actions of other actors on the road, including vehicles and pedestrians, using context and decision-making skills.

For example, when entering a roundabout without any traffic signals for guidance, a driver must judge whether there is a reasonable gap for entry by predicting what other drivers will do and observing nonverbal cues. Often, drivers in vehicles may perform hand gestures to signal that they will give space for entry. In such situations, an autonomous vehicle must be able to interpret signals from surroundings to judge whether there is room for entry.

Another example of a situation that may pose a great challenge for an autonomous vehicle is if a broken-down truck is partially blocking the road. While a human driver would pass such a truck or could roll down the window and further inquire if necessary, an autonomous vehicle will not act in the same manner. If passing the truck requires crossing a double yellow line, the autonomous vehicle will stop until the truck is towed to avoid breaking the law.

Such scenarios that humans consider simple are remarkably difficult for a computer to accomplish without cognitive functions, years of context, and the ability to signal and receive nonverbal cues that humans possess.

The ability for a vehicle to perform such tasks does not come in the form of programmable instructions, but rather requires artificial intelligence, and no AI technologies to date comes close to replicating human intelligence.

For autonomous vehicles to satisfy the demands of driving, such as navigating construction sites, traffic cops, left turns at intersections, and pedestrians, greater precision in electronics and advancements in artificial intelligence are required.

PRECISION OF SOFTWARE AND ELECTRONICS

As Steven Shaldover, pioneer of the ITS and California PATH Program, prompted readers in his article “The Truth About ‘Self- Driving’ Cars,” imagine if your car froze as many times as your laptop or phone.[6] Surely, fewer individuals would be willing to fall asleep in the back seat of an autonomous vehicle if that was the case. Glitches in software that are acceptable for devices like phones and laptops can have detrimental consequences for a computer driving in heavy traffic. Similarly, for AI to assess surroundings and make driving decisions, it requires devices to convey information about the environment and position of the vehicle. Such devices include ultrasonic, radar, LiDAR, and position sensors. Without back-up from human drivers, these devices and software need extreme precision and higher standards than previously demanded to meet safety requirements.

Koopman and Wagner explain in the *SAE Journal of Transportation and Safety* that ADAS systems allow for lower integrity devices since errors can be fixed by human drivers; however, in level 4 and 5 systems, where humans are not expected to take control of the vehicle in response to errors, devices and software must be designed at a higher Automotive Safety Integrity Level (ASIL).[7] Meeting higher standards is not only more difficult to achieve, but also increases the cost and time required for testing as devices must meet lower failure rates with statistical significance.

To satisfy such standards, Koopman and Wagner propose an actuator coupled with a monitor. With this approach, the monitor is designed at high ASIL to detect all faults in the actuator with low false negatives and false positives, while the actuator can be designed to lower standards.[8] If the actuator acts improperly, the monitor will initiate a system shutdown as depicted in Figure 2.

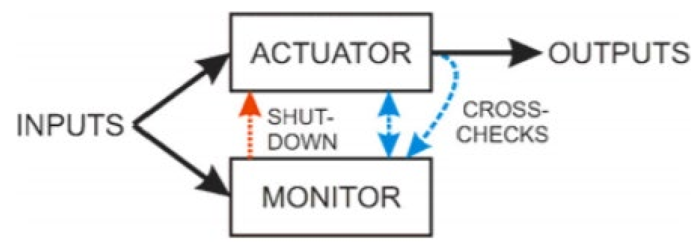

Figure 2. Monitor/actuator pair. This diagram is from Koopman and Wagner.

As seen in Figure 2, the monitor interrupts the actuator output and shuts down the system, acting as a fail-safe once an actuator fault is detected. In practice, the system must remain fail-operational in the event of a fault detection as an actuator shut-down can be dangerous to other drivers on the road who will not anticipate such a response from the autonomous vehicle. As Koopman and Wagner mention, a fail-operational system is challenging to implement and requires additional redundancies.[9]

Other considerations more specific to AV include cybersecurity. Due to the potential for autonomous vehicles to cause damage via collisions, such vehicles are a potential target of attack. Luckily, these threats may not be as significant as feared. Jackson Hickey, vice president of software security firm, Vínsula, reported in an interview with Fagnant and Kockelman that cyber-attacks are more frequently seen in cases regarding espionage and theft of data rather than cases of attack.[10] Performing an attack on autonomous vehicles would be extremely difficult. However, rest assured, cybersecurity will still be of great importance and incur large costs.

Lastly, weather continues to be a great challenge for sensor manufacturers for autonomous vehicles. In a 2019 Design News article, Stewart Sellars, general manager of the Lidar Group for Analog Devices, Inc., admitted “snow is basically an occlusion. It blocks the ability for those sensors [cameras, LiDar, or radar] to get their signal back.”[11] Furthermore, snow on roads block lane markings, creating difficulties to systems such as lane-keeping that depend on these. It is no wonder that Waymo and competitors perform testing in locations such as Arizona with perfect weather conditions and well-marked roads.

Thus, designing systems to meet the higher integrity standards, difficult weather conditions, and security needs of autonomous vehicles creates large costs and challenges for engineers not faced with ADAS systems.

ARTIFICIAL INTELLIGENCE FOR AV

What makes humans suitable drivers is our ability to adapt to new encounters given years of prior context. The environment has countless variables, so every driving experience contains variations. While current technology can succeed in detecting vehicles in blind spots, fast emergency responses, and vehicle-to-vehicle communication, new encounters pose great difficulties to existing systems lacking the ability to contextualize and learn as humans so easily do.

For instance, while an autonomous vehicle may easily detect and respond quickly to an object flying into the road. It does not have the context to identify that the object is in fact cardboard. While a human would recognize the object as benign and proceed normally, an autonomous vehicle may stop in the middle of the road or swerve to avoid the object as it is unable to discriminate between benign and hazardous objects.

Current tests used to measure the capabilities of existing AI technologies reveal that engineers are a long way from creating systems with cognitive functions of humans.

Gary Marcus, director of Uber AI labs, gathered Francesca Rossi, former president of the International Joint Conferences on Artificial Intelligence, and the AI community to develop new tests for artificial intelligence in place of the Turing test, which tested to see if a computer program could trick a judge into believing it was human.

One test, titled the Winograd's Schema Challenge, examined whether systems could identify who or what the pronoun in a sentence referred to based on context. The winner of the Winograd schema was able to answer only 58 percent of the questions correctly[12], indicating the inability of current artificial intelligence systems to comprehend based on context clues.

While Winograd's Schema Challenge demonstrates a lack of ability to interpret context in the case of language, the Physically Embodied Turing Test provides a conclusion even more relevant to autonomous driving. This test asks AI systems to put together furniture from a manual or build objects such as a house from materials and explain the rationale behind each decision. Marcus described the ability to pass this test as "science-fictional."[13]

While an autonomous vehicle will never need to perform such a feat, the "impossibility" of such task emphasizes some of the more relevant challenges to developing level 5 autonomous vehicles. A robot can pick up and move objects with ease. So, if programmed step-by-step instructions of what specific dishes to pick up, how to pick them up, and where to place the different dishes, a robot would be able to unload a dishwasher. However, if the robot was then asked to unload the dishes at another home, it would be an enormous pain. The robot, unable to recognize the new dishes would no longer know what to grasp or how to store the dishes in the new environment.

As demonstrated from the example and results from the new Turing test, without the vast library of experiences that humans have developed over the years and the ability to learn, adapting to changing variables faced in everyday driving poses a serious challenge to autonomous vehicles. While the technology to teach robots to learn like humans are being developed, autonomous vehicle designers have resorted to machine learning and analyzing data to "teach" systems to identify actors on the road and respond appropriately.

To do so, vast quantities of rare encounters in driving, known as edge cases, are shown to systems via virtual simulations. Machines then devise their own rules from the abundance of examples observed, so they will not have to adapt to new situations on the road. However, this approach, known as inductive learning, has several limitations.

The primary limitation is AI 'blind spots,' also referred to as black swans.[14][15] This occurs when special cases are unlearned within a training set due to the proliferation of more frequent data points. For instance, training footage may abundantly show a pedestrian as an individual with two legs, so a machine may fail to label an individual in a wheelchair as a pedestrian.[16] Furthermore, machine definitions and rules learned from training sets are difficult to understand for humans.[17] This illegibility makes it challenging to identify blind spots and correct for them appropriately.

A solution to the problem described above, proposed by Koopman and Wagner, is the operational concept approach, or in simpler terms, restricted applications.[18] This approach limits autonomy to situations containing less variability and thus, fewer edge cases, such as freeways. From here, conditions where autonomous driving is enabled can expand as additional behaviors are learned with more edge cases being simulated and advancements in technology. Recognizing the limitations of inductive learning, CEO of Ford Motor Co., Jim Hackett, acknowledged at the Detroit Economic Club event, "We overestimated the arrival of autonomous vehicles…it's applications will be narrow, what we call geo-fenced, because the problem is so complex."[19] Hackett's statement supports that while level 4 autonomy may be possible in the nearer future, level 5 autonomy is still several decades away.

CHALLENGES TO TESTING

The 2017 safety report from the National Highway Traffic Safety Administration shows that a fatal crash occurs once per 100 million vehicle miles traveled while injury occurs once for every million vehicle miles traveled.[20] Autonomous vehicles must demonstrate that they are at least just as safe.

Completing the billions of hours of physical testing necessary to demonstrate meeting safety requirements with statistical significance would take fleets of autonomous vehicles several hundred years. Furthermore, testing on roads poses a safety hazard to the public. If performed in controlled environments to avoid such safety issues, physical testing would prove extremely costly and time consuming as it would require multiple vehicles and actors to set-up and would be extremely difficult to test all edge cases necessary to deem an autonomous vehicle safe for the road. Unlike testing of standard vehicles that only require a vehicle to meet certain specifications, with autonomous vehicles, these specifications alone are not sufficient. Instead, autonomous vehicles must demonstrate their ability to act appropriately in all edge cases it may encounter on the road.

Thus, researchers have looked to augmented reality to perform testing. MCITY, at the University of Michigan campus in Ann Arbor, is a testing track for autonomous vehicles where computer simulations create virtual vehicles and pedestrians and drop them on the road.[21] The testing vehicle then sees these digital actors through its sensors and responds to the situation accordingly. Researchers plan to create a library of artificial edge cases and run the simulations to test the vehicle's performance and identify failures.[22] Coupled with computers that can run virtual mileage to identify rare edge cases where flaws in vehicle response appear[23], the two methods provide a more feasible alternative for completing testing.

Once again, this solution has its limitations as it is difficult to test for all possible edge cases and identify blind spots when there are several variables encountered in driving environments. Consequently, testers look to restrict the environment to settings with fewer variables and demonstrate safety there first before tackling intersections, construction sites, or snowstorms. Therefore, challenges to testing safety indicate once again that autonomous vehicles, especially level 5 systems, are further away than anticipated.

DANGERS OF LEVEL 3 AUTOMATION

While the arrival of level 4 and 5 systems are delayed due to challenges in meeting higher integrity standards, satisfying all edge cases, and completing testing, level 3 systems are a more immediate possibility; however, does that mean they will be in the hands of consumers soon?

Since level 3 systems require human monitoring and back-up, many researchers fear that such systems will result in more accidents as the attention of drivers drift with less stimulation as computers take control of the vehicle. Thus, several companies such as Waymo have abandoned level 3 systems and are focusing resources towards developing level 4 systems. Despite the additional precision needed for level 4 autonomy, in environments with fewer variables and dangers, such systems pose a more viable option than level 3 autonomous driving systems. We can expect to see level 4 systems appear sooner in restricted environments such as low-speed college campuses, valet parking, and some freeways.

CONCLUSION

Individuals must look beyond bold articles titled "Autonomous vehicles to carry you to your destination in 2020" to see the fine print stating "exclusions apply to paths containing snow, left-turns, emergency vehicles, construction sites, roundabouts, and pedestrians." While several auto manufacturers boasted for years the imminent arrival of level 5 autonomous vehicles, industry leaders are recognizing the challenges of full autonomy and scaling back forecasts.

Although challenges such as the need for increased precision of electronics and software pose some setbacks, the main issue with level 5 autonomy is the shortcomings of current AI technologies. Lacking the years of context that humans have gathered over their lives, as well as the ability to learn and apply context to new situations as humans so easily do, engineers struggle to design autonomous vehicles that demonstrate the adaptability that make humans good drivers. Given the high variability of driving, finding sufficient methods around teaching a computer cognitive functions have proven difficult.

Current inductive learning methods pose several limitations given the prevalence of blind spots and the difficulty of identifying these blind spots. Thus, designing level 5 autonomous vehicles that respond appropriately to all edge cases imaginable is an impossible task.

Not only is designing such a vehicle extremely difficult, testing the safety of the vehicle to ensure it meets all standards and performs as well as human drivers is another great challenge. Methods for testing, including augmented reality environments and computers running virtual mileage, rely on a library of artificial edge cases that once again runs into the issue of identifying and correcting all blind spots.

With great difficulties in designing and testing level five systems, fully autonomous vehicles are several decades away. Although the technology for level 3 systems are more present, the dangers of these systems may outweigh the benefits and have caused many companies to move away from implementation. Thus, aside from Tesla which hopes to achieve full autonomy from improvements to its level 3 system, we are likely to see level 4 systems from competitors arrive next to the market. Although less impressive than flashy headlines describing robo-taxis, level 4 autonomy still promises benefits to consumers. Meanwhile, the advancements ahead in AI and electronics to make fully autonomous vehicles a possibility are exciting and stand to benefit society in several other facets.

[1] Murray, Charles. “Automakers Are Rethinking the Timetable for Fully Autonomous Cars.” *Design News*, 6 June 2019, https://www.designnews.com/electronics-test/automakers-are-rethinking-timetable-fully-autonomous-cars/93993798360804.
[2] Funk, Jeffrey. “What’s Behind Technological Hype?” *Issues in Science and Technology* 36, no. 1 (Fall 2019): 36–42.
[3] Rowlatt, Justin. “Why You Have (Probably) Already Bought Your Last Car.” *BBC News*, 10 Oct. 2018, https://www.bbc.com/news/business-45786690.
[4] “Automated Vehicles for Safety.” *NHTSA*, 12 Aug. 2019, https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety.
[5] Yoshida, Junko. “Car OEMs See No Easy ADAS-to-AV Path.” *EE Times*, 26 Nov. 2019, https://www.eetimes.com/car-oems-see-no-easy-adas-to-av-path/#.
[6] Shaldover, Steven E. “The Truth about ‘Self-Driving’ Cars.” *Scientific American*, Dec. 2016, https://www.scientificamerican.com/article/the-truth-about-ldquo-self-driving-rdquo-cars/.
[7] Koopman, Philip, and Michael Wagner. “Challenges in Autonomous Vehicle Testing and Validation.” *SAE International Journal of Transportation Safety*, vol. 4, no. 1, 2016, pp. 15–24. *JSTOR*, www.jstor.org/stable/26167741.
[8] Ibid.
[9] Ibid.
[10] Fagnant, Daniel J., and Kara Kockelman. “Preparing a Nation for Autonomous Vehicles: Opportunities, Barriers and Policy Recommendations.” *Transportation Research Part A: Policy and Practice*, vol. 77, 2015, pp. 167–181., doi:10.1016/j.tra.2015.04.003.
[11] Murray, Charles. “Automakers Are Rethinking the Timetable for Fully Autonomous Cars.” *Design News*, 6 June 2019, https://www.designnews.com/electronics-test/automakers-are-rethinking-timetable-fully-autonomous-cars/93993798360804.
[12] Marcus, Gary. “The Search for a New Test of Artificial Intelligence.” *Scientific American*, Mar. 2017, https://www.scientificamerican.com/article/the-search-for-a-new-test-of-artificial-intelligence/.
[13] Ibid.
[14] Massachusetts Institute of Technology. "Self-driving cars, robots: Identifying AI 'blind spots'." ScienceDaily. ScienceDaily, 25 January 2019, www.sciencedaily.com/releases/2019/01/190125094230.ht m.
[15] Koopman, Philip, and Michael Wagner. “Challenges in Autonomous Vehicle Testing and Validation.” *SAE International Journal of Transportation Safety*, vol. 4, no. 1, 2016, pp. 15–24. *JSTOR*, www.jstor.org/stable/26167741.
[16] Ibid.
[17] Ibid.
[18] Ibid.
[19] Murray, Charles. “Automakers Are Rethinking the Timetable for Fully Autonomous Cars.” *Design News*, 6 June 2019, https://www.designnews.com/electronics-test/automakers-are-rethinking-timetable-fully-autonomous-cars/93993798360804.
[20] NHTSA. *Traffic Safety Facts 2017: A Compilation of Motor Vehicle Crash Data*. US Department of Transportation, 2019.
[21] Feng, Yiheng, and Henry X Liu. “Self-Driving Cars Learn About Road Hazards Through Augmented Reality.” *IEEE Spectrum*, 20 Nov. 2019, https://spectrum.ieee.org/transportation/self-driving/selfdriving-cars-learn-about-road-hazards-through-augmented-reality.
[22] Ibid.
[23] “Virtual-Based Safety Testing for Self-Driving Cars from NVIDIA DRIVE Constellation.” *NVIDIA*, https://www.nvidia.com/en-us/self-driving-cars/drive-constellation/.