Statistics, AI, and Autonomous Vehicles

August 29, 2023

David Banks, Duke University, banks@stat.duke.edu
Yen-Chun Liu, Duke University, yenchu.liu@duke.edu

1 Context

Yogi Berra said "It's hard to make predictions, especially about the future." But making quantified predictions is something statisticians are supposed to do pretty well.

In this case, we want to examine the future of statistics in AI for autonomous vehicles (AVs). The exact service that statisticians eventually provide will depend upon how the interplay of technology, law, economics, and social adoption of AVs unfolds. Certainly we shall be called upon to do risk analyses, probably under many different scenarios (e.g., for a mixed fleet, just on highways, and in bad weather). We shall surely make estimates of the impact of AVs on the environment and the economy. We are likely to play a role in insurance, and possibly regulation. And we may be asked to develop statistical procedures to assess the quality and capability of the AI software and the completeness of the AI training.

AVs could be transformational. They have the potential to significantly reduce pollution, save lives, and provide other important benefits. They are one of the few technologies with the potential to meaningfully address climate change.

The greatest benefits would accrue if all vehicles on the road were AVs that were networked together. There would be little need to brake, which is wasteful of fuel—cars could adjust their speeds to interweave without stopping and restarting at stoplights. Platooning would also save fuel. And, if AVs can become sufficiently safe, it would no longer be necessary to carry around 1.5 tons of steel for protection. One could make car bodies out of canvas, saving yet more fuel.

In terms of safety, AVs are never tired or distracted, and they have better sensors than human-piloted vehicles. If networked, the AVs can share information about safety conditions such as a herd of deer near a road or a pothole hidden beneath a puddle.

According to the National Safety Council, of the 42,939 crash fatalities in 2021 (Stewart, 2023), about 20,000 involved multi-car collisions (National Safety Council, 2022). If only networked AVs were on the road, this number would presumably drop to nearly zero. Among the 9,026 single-vehicle fatal crashes in 2021, 59% of drivers had blood alcohol levels above 0.08. Not all of these fatalities are due to alcohol use, but there is much evidence that alcohol is a dominant factor in vehicle deaths (National Highway Transportation Safety Administration, 2021). Networked AVs would also eliminate deaths and injuries caused by inexperienced teen drivers and elderly drivers with diminished capability.

Modern work using software to pilot robotic vehicles began in 1984, when William "Red" Whittaker launched the NavLab project. In 1995, NavLab 5, a small truck carrying a computer, drove from Pittsburgh to San Diego, mostly without human guidance (Jochem et al., 1995). Another roboticist involved with that team at Carnegie Mellon University was Sebastian Thrun, who played a leading role in Google's initial effort to build AVs, which has now been spun off as Waymo. See Burns and Shulgan (2018) for a history of the people and companies working in this area.

There are six levels of vehicle automation:

Level 0. No automation. The human has only standard assistance (mirrors, rear-view cameras).

Level 1. Software assistance. Examples include adaptive cruise control and lane keep assist. This level became widely available after 2018.

Level 2. Partial automation. The driver must be hands-on and ready to take control, but the car controls speed and holds its lane. The Tesla Autopilot is an example.

Level 3. Conditional automation. Hands are off the wheel, but the driver must still be ready to control. It is intended for limited access highways and good driving conditions. Many automobile companies are experimenting with this level.

Level 4. High automation. Driver can sleep after inputting destination. Waymo is testing these kinds of AVs. The car must stay on traditional roads.

Level 5. Full automation. This level enables navigation of non-traditional roads and is years away.

From a statistical risk analysis standpoint, we are beginning to acquire significant amounts of data on Levels 3 and 4.

Waymo is essentially the only company operating Level 4 vehicles. They report that, as of January 2023, they had driven 1 million AV miles (The Waymo Team, 2023). That news release also indicates that there have been no injuries, that there were 18 minor contact events, and of those ten were a human driver hitting a stationary Waymo AV. They assert that in every vehicle-to-vehicle collision, the human operator violated road rules or drove dangerously.

It is difficult to compare the Waymo record to human performance. The number of miles driven is too small to warrant examination based on fatalities, and injuries are difficult to define. Nonetheless, the Bureau of Transportation Statistics (2021) reports that in 2020, there were 2,250,000 roadway injuries and that people in the U.S. drove a total of 2.9 trillion miles. One would have expected 0.78 Waymo injuries if Level 4 AVs are identical to human drivers, and there were none. There is insufficient evidence to conclude that Level 4 vehicles are safer than human drivers, but it strongly suggests that Level 4 AVs are not worse than humans. Note that Waymo operates mostly in the metropolitan areas of Phoenix, AZ, and San Francisco, CA (Bonk, 2023), and injury rates tend to be higher in urban areas

Obviously, statisticians in the Bureau of Transportation Statistics and the National Highway Transportation Safety Administration (NHTSA) are better placed than we are to do risk analyses that control for driving conditions and the kind of accidents that occur. This is clearly an important role for statisticians, and AV safety should be continually assessed.

Since July of 2021, NHTSA has required AV manufacturers to report crash data (National Public Radio, 2022). As of January 15, 2023, NHTSA says carmakers have submitted 419 AV crash reports. Of these, 263 of those accidents have been in Level 2 vehicles, with 156 involving Level 3 or higher AVs (National Highway Transportation Safety Administration, 2023). NHTSA further reports 18 fatalities, all with Level 2 cars. No carmaker has reported a fatality with a Level 3 or higher AV, but it should be noted that there are far fewer of these on the road. California reports that Level 3 AVs drove a total of 10.4 million miles between January 2021 and December 2022 (Office of Public Affairs, State of California, 2023), but that is too small to warrant a risk assessment, since, in 2020, the fatality rate per 100 million miles driven is only 1.34 (National Highway Transportation Safety Administration, 2022).

The NHTSA report has 19 accidents for which the injury level is listed as “unknown”. Also, NHTSA does not require carmakers to indicate whether the AV was at fault in the crash, nor whether the accident was caused by user error or a problem with the AI. Nonetheless, the statisticians at NHTSA are uniquely situated undertake fine-grained risk analyses which, among other purposes, should flag common failure modes. Also, it is obvious that any current risk analyses pertain to the mixed fleet situation, and there is reason to think that the greatest safety benefits will only accrue when all vehicles on the road are networked AVs.

2 Scenarios

There are several ways in which the future of AVs might develop. One is the situation we have now, in which traffic is a mix of vehicle levels. The available data hint that, on average, AVs are somewhat safer than human-driven cars. This scenario will evolve as car manufacturers and statisticians learn more about the conditions under which a specific Level of AV performs well or poorly. We can imagine that one day a person might unlock their Level 2 AV and the car will say ”You will have to drive yourself today. There is snow on the road and the AI feels it is not able to operate the automobile safely.”

Inevitably, this scenario will lead to new regulations, legal decisions, and insurance policies, among many other changes. The regulations would probably prescribe how software updates should be made securely, in a world where international cyberattacks have become common (Greenberg, 2018). Similarly, they would govern inspection frequency, the precise capability of different levels of AV, use of seatbelts, where AVs could be driven (perhaps just on highways, for some levels), and so forth. Legal decisions would determine liability in crashes, which in turn would drive new forms of insurance policies. As in all systems that involve industry and government, we should expect the future to be shaped by path dependence (Bebchuk and Roe, 1999), which implies that change will be incremental and shaped by previous choices, rather than a revolutionary attempt to optimize outcomes.

A second scenario, not entirely incompatible with the first, is that a significant force for AV adoption will come from the trucking industry and/or the ride-share industry. Both will profit if driving can be automated (with additional benefits from all-electric fleets). Perhaps a bargain will be struck in which one lane of the Interstate highway system will be dedicated to AV trucks. When an AV truck leaves the highway, one can imagine a regulation requiring a human operator to take full or partial control, as air force pilots direct international drones from domestic bases. Similarly, Waymo is operating a ride-hailing system in Phoenix and San Francisco, and will surely lobby to speed the transition to an AV world. And change may happen in tandem—Uber is building a freight fleet that uses Waymo’s self-driving trucks (McFarland, 2022).

A third aspect of an AV future is that probably many vehicles will look very different than transportation built to move humans. During the pandemic, there has been a huge uptick in online purchases (Koch et al., 2020), which leads to deliveries that encounter the “last mile” problem (Song et al., 2009). There are economically efficient ways to move goods from the site of manufacture to the depots in the cities in which purchasers live, but then fleets of panel trucks are needed to carry the items to the home of the purchaser. In an AV future, these panel trucks might be replaced by something that looks more like an electrically powered grocery store shopping cart.

It is entirely possible that the world will move away from the concept of car ownership. Instead, people might gravitate towards the Waymo/Uber ride-hailing system (Bellan, 2023). Among many advantages, a family of six could summon a large van when traveling together, or a fuel-efficient single-seater when only one person is going, or a luxury vehicle when traveling a long distance, or a truck when moving house. Additionally, a ride-hailing system better lends itself to use of electric vehicles, since the AV can take itself off the road to recharge, which is good for the environment.

3 Training an AV

All AVs use AI in the form of deep neural networks. Training a deep neural network requires lots of data, an appropriate architecture, and an optimization algorithm. The architecture should be deep, and in practice it is often modular with subnetworks trained to perform specific recognition or prediction tasks. The optimization algorithms are proprietary, but probably quite complex and specific to given subproblems. One of the limiting factors in training neural networks is that they rely upon human beings to categorize unrecognized images, such as a horse on a road or a heavily weathered stop sign.

There are many constraints. The latency must be short in order for the AI to react quickly to changing traffic conditions. Neural network code must be simple enough to store on-board (nevertheless, it is quite complex). The implementations evolve over time, and are highly proprietary. Statistical thinking arises in the training of deep neural networks, but it is not central, and so our review of this aspect of AVs will be brief.

Initially, the goal was to have an AV that could hold its lane and follow a car in front of it at at fixed distance. But, inevitably, more complex behavior is wanted. In the Tesla AV, the architecture is complex. There are multiple camera sensors, whose raw images are processed by a rectification layer in the deep network, to correct for miscalibrated cameras. These images are then sent to a residual network which processes them into features at different channels and scales—features might correspond to stop signs, lane markings, and other vehicles. Features are then fused into multi-scale information using a module that represents it in a vector space. The output space is sent to a spatio-temporal queue which is processed by a recurrent neural network. Its output is fed into the branching structure of a component called the hydranet, which sends the processed data to many different modules that handle different prediction task. For much more detail, including a link to a YouTube video on the topic by Andrej Karpathy, Tesla's director of AI and

Autopilot Vision, see Van Laer (2022).

Training an AV AI can be gamed. Eykholt et al. (2018) showed how applying a post-it note to a stop sign can fool an AI into thinking that the stop sign is a billboard. Similarly, if the training data has been "poisoned", the AI may be misled (Modas et al., 2019). Altering just a few pixels in an image can cause problems.

Regarding training, it is worth noting that Tesla has driven far more miles than Waymo or any other AV system. As of July, 2022, it had driven 35 million miles (Lambert, 2022). As of 2019 it had a fleet of around 5,000 vehicles and drove as much in a day as Waymo has driven to date (Bouchard, 2019). When a Tesla AI encounters a new situation, or "sees" something it doesn't recognize, it records it to add to the training database. This means that it has more data, and thus has the potential to jump from Level 3 to Level 4. Having more data and rarer data is a critical competitive advantage when training a deep neural network.

4 Impact

We believe that AVs will have major economic impact. They can substantially lower the cost of production (less expensive supply lines and lower costs for distributing goods). This will happen because of greater fuel efficiency, and because drivers are no longer needed to pilot trucks to destinations. Such change would cause significant dislocation in employment, and an ethical society should begin now to plan how to ameliorate the impact on people who lose jobs in exchange for the benefit of everyone.

We also think that AVs will have significant social impact. It could move us away from privately owned cars towards a more ubiquitous ride hailing system. Perhaps children could be delivered to school by an AV, rather than being dropped off by parents. And the elderly could stay in their homes longer, even after no longer being able to drive (National Institute of Health, 2022). Before COVID-19, the average time an American spent commuting was 55.2 minutes per day (United States Census Bureau, 2020). AVs would allow people to read, work or sleep, giving them an extra hour of discretionary time each day.

But the economic and social impact of AVs is dwarfed by their potential environmental impact. The world's population is increasing. According to the U.S. Census Bureau, there are 7.97 billion people now (United States Census Bureau, 2023), and by 2050, that number will be between 9.5 billion and 10 billion. Figure 1 shows the estimated total, and also estimates that are broken out be region.

According to researchers at the Australian Academy of Science, the carrying capacity of the planet, with anything like a middle-class American lifestyle, is on the order of 2

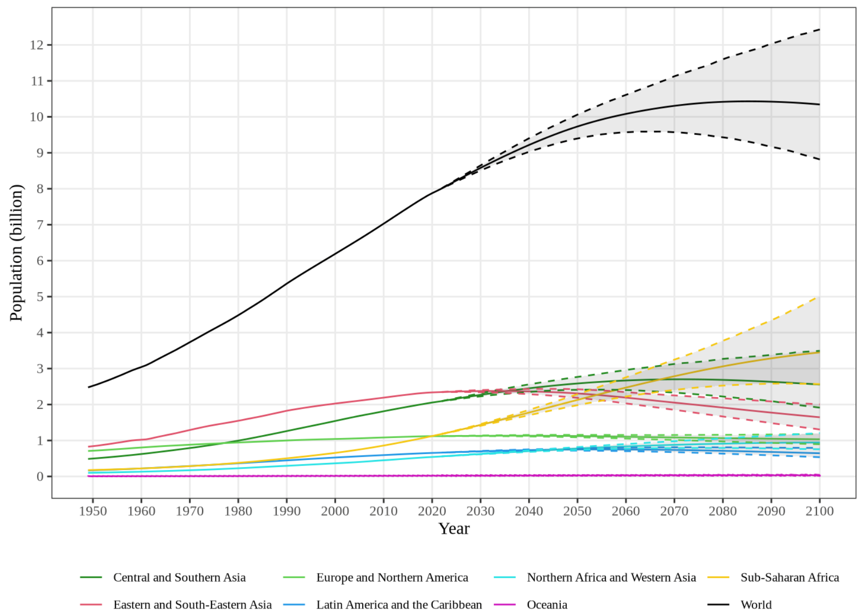

Figure 1: The graph shows the recent and projected human population of the world, and also breaks that information out by region.

billion (Dovers and Butler, 2017). There is no technology on the horizon that can convert a ton sand into a ton of food and use very little energy to do it. But that is the scale of the problems we face.

Climate forecasting is less accurate than demography, but scientists at NASA estimate that by 2050, parts of South Asia, the Persian Gulf, and the Red Sea will be too hot to support human habitation Buis (2022). And by 2070, the same will be true for parts of Brazil, eastern China, and southeast Asia. Geopolitically, there will be consequences. The people of Bangladesh (who also face problems from the rising sea level) will have to move north to Pakistan or Afghanistan—a relocation that seems fraught with political and military peril. Similarly, the political governance in Syria and Lebanon is precarious. It is easy to imagine that climate pressure, say a week of 120 degree temperatures, could lead to collapse and further waves of refugees.

The only way for people in such regions to survive is through access to potable water and a high level of technology. Most of the nations in those regions do not have such capacity now, and may not acquire the necessary resources in time. But what seems certain is that there is no government in the world that has the political will to leave coal in the ground unburned when it is 120 degrees outside and the nation needs to run a desalinization plant. That sets up a vicious cycle, in which more carbon is produced, causing higher temperatures, which in turn leads to more carbon being released.

Which brings us back to AVs. They are one of the few technologies that could actually have a meaningful impact on our future's carbon footprint. Transportation accounts for 29% of U.S. greenhouse gas emissions—it is the largest single component of our pollution budget (Environmental Protection Agency, 2023a). In the United States, transportation accounts for 41.7% of CO_2 emissions (77% if one excludes wildfires) (Environmental Protection Agency, 2023b, National Tier 1 CAPS spreadsheet). Networked, electric AVs could dramatically mitigate global warming, while simultaneously impproving our quality of life, socially and economically.

5 Conclusions

AVs have enormous potential impact, and forward planning can achieve important social and economic benefits, as well as essential environmental benefits. AVs depend upon AIs, which are trained upon vast quantities of data. Statistical methods are key to many aspects of AV implementation. Risk analysis is obvious, but our community can and should play a much larger role.

Examination of the references in this paper shows heavy citation of work by statisticians at the U.S. Census Bureau, the Bureau of Transportation Statistics, the National Highway Transportation Safety Administration, the Environmental Protection Agency, and many other sources of official statistics. That work is the foundation of this paper, and demonstrates the ubiquity of statistics in developing the AV application of AI.

References

Lucian Arye Bebchuk and Mark J Roe. A theory of path dependence in corporate ownership and governance. *Stan. L. Rev.*, 52:127, 1999.

R. Bellan. Waymo's Self-Driving Cars will be Available on Uber's App, Starting in Phoenix, 2023. Available at `https://techcrunch.com/2023/05/23/waymos-self-driving-cars-will-be-available-on-ubers-app-starting-in-phoenix/`, cited on 6/7/2023.

L. Bonk. Waymo Doubles Service Area for Its Fleet of Robo-taxis, 2023. Available at `https://www.engadget.com/waymo-doubles-service-area-for-its-fleet-of-robo-taxis-163353714.html`, cited on 6/7/2023.

Y. Bouchard. Tesla's Deep Learning at Scale: Using Billions of Miles to Train Neural Networks, 2019. Available at https://towardsdatascience.com/teslas-deep-learning-at-scale-7eed85b235d3, cited on 6/7/2023.

A. Buis. Too Hot to Handle: How Climate Change May Make Some Places Too Hot to Live, 2022. Available at https://climate.nasa.gov/explore/ask-nasa-climate/3151/too-hot-to-handle-how-climate-change-may-make-some-places-too-hot-to-live/#:~:text=Future%20Hot%20Spots&text=But%20climate%20models%20tell%20us,Asia%2C%20and%20Brazil%20by%202070., cited on 6/11/2023.

Bureau of Transportation Statistics. Injured Persons by Transportation Mode, 2021. Available at https://www.bts.gov/content/injured-persons-transportation-mode, cited on 6/12/2023.

Lawrence D Burns and Christopher Shulgan. *Autonomy: The Quest to Build the Driverless Car—And How It Will Reshape Our World*. HarperCollins, 2018.

S. Dovers and C. Butler. How Many People Can Earth Actually Support?, 2017. Available at https://www.science.org.au/curious/earth-environment/population-environment, cited on 6/11/2023.

Environmental Protection Agency. Carbon Pollution from Transportation, 2023a. Available at https://www.epa.gov/transportation-air-pollution-and-climate-change/carbon-pollution-transportation#:~:text=%E2%80%8BGreenhouse%20gas%20(GHG)%20emissions,contributor%20of%20U.S.%20GHG%20emissions., cited on 6/11/2023.

Environmental Protection Agency. Air Pollutant Emissions Trends Data, 2023b. Available at https://www.epa.gov/air-emissions-inventories/air-pollutant-emissions-trends-data, cited on 6/12/2023.

Kevin Eykholt, Ivan Evtimov, Earlence Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1625–1634, 2018.

Andy Greenberg. The untold story of notpetya, the most devastating cyberattack in history. *Wired, August*, 22, 2018.

Todd Jochem, Dean Pomerleau, Bala Kumar, and Jeremy Armstrong. Pans: A portable navigation platform. In *Proceedings of the Intelligent Vehicles' 95. Symposium*, pages 107–112. IEEE, 1995.

Julia Koch, Britta Frommeyer, and Gerhard Schewe. Online shopping motives during the covid-19 pandemic—lessons from the crisis. *Sustainability*, 12(24):10247, 2020.

F. Lambert. Tesla Surpasses 35 Million Miles Driven on Full Self-Driving Beta, and Pace Is Ramping Up, 2022. Available at `https://electrek.co/2022/07/20/tesla-surpasses-35-million-miles-driven-full-self-driving-beta/`, cited on 6/7/2023.

M. McFarland. Uber and Waymo Fought in Court over Self-Driving Trucks. Now They're Teaming Up, 2022. Available at `https://www.cnn.com/2022/06/07/tech/uber-waymo-self-driving-deal/index.html`, cited on 6/7/2023.

Apostolos Modas, Seyed-Mohsen Moosavi-Dezfooli, and Pascal Frossard. Sparsefool: a few pixels make a big difference. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9087–9096, 2019.

National Highway Transportation Safety Administration. Drunk Driving., 2021. Available at `https://www.nhtsa.gov/risky-driving/drunk-driving`, cited on 6/6/2023.

National Highway Transportation Safety Administration. NHTSA Releases 2020 Traffic Crash Data, 2022. Available at `https://www.nhtsa.gov/press-releases/2020-traffic-crash-data-fatalities#:~:text=While%20the%20number%20of%20crashes%20and%20traffic%20injuries,increase%20from%202019%20and%20the%20highest%20since%202007.`, cited on 6/7/2023.

National Highway Transportation Safety Administration. Summary Report: Standing General Order on Crash Reporting for Level 2 Advanced Driver Assistance Systems, 2023. Available at `https://www.nhtsa.gov/sites/nhtsa.gov/files/2022-06/ADAS-L2-SGO-Report-June-2022.pdf`, cited on 6/7/2023.

National Institute of Health. Safe Driving for Older Adults, 2022. Available at `https://www.nia.nih.gov/health/older-drivers#:~:text=As%20you%20age%2C%20your%20joints,wheel%20quickly%2C%20or%20brake%20safely.`, cited on 6/12/2023.

National Public Radio. Nearly 400 Car Crashes in 11 Months Involved Automated Tech, Companies Tell Regulators., 2022. Available at https://www.npr.org/2022/06/15/1105252793/nearly-400-car-crashes-in-11-months-involved-automated-tech-companies-tell-regul, cited on 6/6/2023.

National Safety Council. Injury Facts, 2022. Available at https://injuryfacts.nsc.org/motor-vehicle/overview/type-of-crash/, cited on 5/20/2023.

Office of Public Affairs, State of California. Autonomous Vehicle Permit Holders Report 5.7 Million Test Miles in California., 2023. Available at https://www.dmv.ca.gov/portal/news-and-media/autonomous-vehicle-permit-holders-report-5-7-million-test-miles-in-california/, cited on 5/31/2023.

Liying Song, Tom Cherrett, Fraser McLeod, and Wei Guan. Addressing the last mile problem: transport impacts of collection and delivery points. *Transportation research record*, 2097(1):9–18, 2009.

T. Stewart. An Overview of Motor Vehicle Traffic Crashes in 2021. 2023. Report No. DOT HS 813 435). National Highway Traffic Safety Administration.

The Waymo Team. First Million Rider-Only Miles: How the Waymo Driver is Improving Road Safety., 2023. Available at https://blog.waymo.com/2023/02/first-million-rider-only-miles-how.html, cited on 5/31/2023.

United States Census Bureau. Travel Time to Work in the United States: 2019, 2020. Available at https://www.census.gov/newsroom/press-releases/2021/one-way-travel-time-to-work-rises.html, cited on 6/12/2023.

United States Census Bureau. U.S. and World Population Clock, 2023. Available at https://www.census.gov/popclock/, cited on 5/25/2023.

I. Van Laer. Tesla's Self-Driving Algorithm Explained, 2022. Available at https://towardsai.net/p/l/teslas-self-driving-algorithm-explained#:~:text=Each%20specific%20task%20has%20its,layout%20is%20called%20a%20HydraNet., cited on 6/7/2023.