UNIVERSITY OF SAN FRANCISCO

CHANGE THE WORLD FROM HERE

When Noone is Driving

Navigating Accountability via Cruise's Driverless Vehicles

Themes:

AI Regulation
Transparency & Explainability
Social & Economic Impact

Prerequisites:

- None for the Case Study section
- Basic Python for the Technical Exercises, with some knowledge of classes

Owner:

Center for AI and Data Ethics at University of San Francisco

Author(s):

Hadley Dixon and Robert Clements

License:

Citation:

Dixon, Hadley and Clements, Robert. (2024). When Noone is Driving: Navigating Accountability via Cruise's Driverless Vehicles.

Objective:

The purpose of this case study is to explore the many questions around ownership and accountability when it comes to artificial intelligence used to create autonomous vehicles.

Instructions:

1. Read through the case study individually and then answer the discussion questions as a group, or in small groups.
2. Technical audiences: individually complete the exercises.
3. Take a ride in a driverless vehicle…or not.

Case Study:

Autonomous vehicles have gained traction in the last 20 years, particularly in technology-dense areas like Silicon Valley. Rapid developments in artificial intelligence (AI) and machine learning have made it possible for vehicles to "learn" how to drive by using self-improving algorithms, and now they can make real-time decisions and, in some conditions, can drive similar to human drivers. In 2015, this technology was popularized as Tesla launched its Autopilot feature, which was designed to navigate the roads with minimal driver involvement. Since then, many companies, including Waymo and Cruise, have joined Tesla in a race to perfect self-driving technology. Waymo deployed its fully autonomous vehicle on the streets of Phoenix in 2020, followed by deployment in other large cities like San Francisco. Cruise followed close behind, backed by both General Motors and Honda, with extensive user testing in the immediate San Francisco area. In 2022, Cruise launched commercial autonomous rides, offering both ride-hailing services to the public for a fee and free rides to local college students, over the age of 18. With increased visibility of autonomous vehicles, Cruise, Waymo and other developers seek to reduce reliance on traditional vehicles with a safe, convenient and accessible AI-powered replacement.

However, not all California residents are thrilled about these developments. Public concerns are diverse, the majority of which center around accountability. The Washington Post reports on one highly controversial incident in San Francisco regarding a crash that sent a woman to the hospital with life-threatening injuries (Thadani, 2023). In October 2023, on the corner of 5th and Market St., at around 9:30 PM, a 911 call reported a vehicle hitting a woman walking across the street. When first responders arrived at the scene of the crime, they found a woman pinned underneath one of Cruise's driverless vehicles, with her leg sticking out from underneath the vehicle's left rear wheel. There were no other vehicles around, as well as no passengers inside the Cruise vehicle. Using Cruise's voice controls, first responders were able to contact an operator and disable the vehicle.

The day following the accident, Cruise released a statement, stating that their vehicle is not at fault. The company's spokesperson said, "the initial impact was severe and launched the pedestrian directly in front of the AV." A recording from the self-driving vehicle's dash cam revealed a regular vehicle colliding with the woman and throwing

her into the path of Cruise's autonomous vehicle before the initial vehicle flees the scene. In response, Cruise's vehicle braked aggressively but was unable to stop completely before rolling over the woman and coming to a final halt. Dash cam video also shows both vehicles having a green light, and the pedestrian attempting to cross the street into oncoming traffic. In the weeks to follow, Cruise worked closely with law enforcement to identify the responsible driver. San Francisco police confirmed the basic details offered by Cruise but did not call this accident a hit-and-run.

Discussion Questions:

1. Should we test experimental technology on actual city streets, even if their presence impacts important agencies such as emergency services? How about rural and suburban areas?
2. Are companies transparent enough about the risks, and does it matter to our political leaders?
3. Can, and should, developers ensure that their algorithms prioritize the safety and well-being of pedestrians and other road users over the autonomous vehicle's passengers or passengers of other vehicles on the road?
4. In this incident, is the driver of the first vehicle that struck the pedestrian accountable for the accident, or is it the pedestrian, or some combination of the vehicles and the pedestrian? Does the Cruise vehicle deserve any of the blame?
5. The Cruise vehicle in this case chose to brake hard, and then essentially shut down. Why might it have made this decision? How can we find out?

After several months of on-the-street testing, California authorized Waymo and Cruise to offer paid robotaxi service in San Francisco. This new rideshare service quickly triggered several investigations of blame regarding recent accidents. On one hand, you have developers who insist that the news sensationalizes autonomous vehicles and that the presence of Waymo and Cruise, and eventually broadband autonomous driving, will make streets safer. On the other hand, prominent politicians like House Speaker Nancy Pelosi claim that "AVs [autonomous vehicles] are not merely an inconvenience; they endanger the lives of passengers, other drivers, pedestrians, and even individuals with no proximity to the vehicles who require emergency services." Pelosi's claims are echoed by emergency services throughout San Francisco. SF fire and police departments blame autonomous vehicles for blocking ambulances from responding to emergency scenes. The fire department also documented driverless vehicles rolling over fire hoses used to put out fires, and have tallied 55 incidents where self-driving vehicles have gotten in the way of rescue operations (Kerr, 2023). Such calls have resulted in both Cruise and Waymo cutting their fleets by 50%, followed by a recall of all Cruise vehicles by the San Francisco DMV in November 2023.

In situations such as these, with no drivers or passengers, as well as no public witnesses on the scene, questions about accountability are particularly difficult to answer. For partially autonomous vehicles, which still involve human control, assigning liability depends on the actions leading up to the collision. For fully autonomous vehicles, like the ones operated by Cruise and Waymo, blame may be shared between

manufacturers, the service center, the vehicle owner, other human-operated vehicles, among others. For example, negligence liability might fall on the vehicle owner for failing to implement a software update from the manufacturer. In other cases, like in this scenario, Cruise's shutdown feature that was deployed while the woman was underneath the vehicle is called into question. For developers, this controversy is a matter of transparency and accountability. Many claim that individuals have the right to understand how AI systems are making decisions that affect them. The black box nature of models used to power autonomous vehicles make it hard to explain decision-making. Machine learning algorithms, sensors, and real-time data introduces challenges in discerning whether an accident resulted from a flaw in the vehicle's programming, sensor inaccuracies, or unforeseeable external factors. Furthermore, the same sensor data that could trigger an accident could also inform liability decisions. Because AV systems track, monitor, and measure the surroundings with a range of sensors, they can be a crucial tool in determining fault. It is a double-edged sword. Parties may have disproportional control over sensor data, so the issue of transparency becomes relevant again. It is important to hold developers and users of AV systems accountable for their actions, while balancing improvement of these technologies.

It is important to recognize that the advent of autonomous vehicles is only one symptom of the gradual automation of the economy that's been unfolding for the last 15 years. "Gig work" is characterized as freelancing work by short-term, independent contractors. It gained prominence primarily through rideshare apps like Uber, Lyft, and food delivery services like DoorDash and Instacart. This new "gig economy" has grown in popularity over the last decade and its economic impact is expansive, fundamentally reshaping the way people work and access services. The appeal of this model offers flexibility in working hours and some gig workers see it as an attractive alternative to traditional employment. However, gig workers often lack the stability and protections afforded by traditional employment, such as healthcare, job security, and minimum wage guarantees. As a result, gig workers are vulnerable to labor exploitation by large corporations, who are not legally required to provide essential benefits like paid leave or retirement plans. Questions of inequality, worker exploitation, and insecurity are heightened as systematic autonomization of the economy continues.

The impact of automation is far-reaching, and has begun reshaping many existing industries and employment dynamics. Proponents argue that automation improves efficiency by streamlining processes, reducing human error, and increasing production. However, this economic growth comes at a cost that puts the burden squarely on the shoulders of workers. As industries begin incorporating automation in everyday processes, many workers fear for their jobs and are forced to confront the possibility of unemployment and economic dislocation in the near future. Automation drives a shift in the types of skills demanded in the labor market, which for many "low-skilled" workers, means the eventual obsolescence of repetitive or manual labor. Focusing specifically on the impact on the gig economy, we see how automation magnifies ethical concerns and brings a new perspective on the future of existing gig work.

Development of autonomous vehicles naturally brings the prospect of completely self-driving rideshare services to the forefront for companies like Uber and Lyft. These self-driving cars, which promise efficiency, safety, and objective decision-making, offer a

cost-effective incentive for corporations to significantly reduce their investment in human capital, opting instead for an entirely autonomous fleet. Notably, nearly half of freelancers consider their gig work to be a long-term career move rather than a short-term opportunity (Sabatier, 2024). This shift could impact the livelihoods of thousands of gig workers who rely on driving for their income. Additionally, drivers have voiced concerns regarding privacy and surveillance, as gig workers are subject to regular monitoring of performance, ratings, and behavior, as well as 24/7 location tracking. This surveillance poses significant autonomy issues, as workers feel pressured to meet arbitrary performance metrics to maintain their livelihood, which is not guaranteed by any employee contract.

Discussion Questions:

6. With the transition to autonomous vehicles, how will revenue generated from transportation services now flow into different sectors of the economy (tech companies, drivers, etc.)?
7. In the wake of automation, what measures can be implemented to ensure equitable access to transportation services that many workers have begun to financially rely upon?
8. What role, if any, should the government play in regulating the deployment of autonomous vehicles to mitigate the impact on gig workers?
9. Outside of crashes, and the economy, what are ethical considerations need to be discussed?

References:

Geneva Internet Platform (n.d.). *The rise of autonomous vehicles*. Digital Watch. https://dig.watch/trends/rise-autonomous-vehicles

Haghani, M. (2024, January 29). *What moral principles should self-driving car algorithms follow in situations where harm is unavoidable?* ABC Religion & Ethics. https://www.abc.net.au/religion/what-moral-principles-should-guide-self-driving-ai-cars/103403186

University of Auckland (2023, December 14). *Ethics on autopilot: The safety dilemma of self-driving cars*. Tech Explore. https://techxplore.com/news/2023-12-ethics-autopilot-safety-dilemma-self-driving.html

Jurdak, R., & Kenhere, A. (2018, March 21). *Who's to Blame When a Driverless Car Has an Accident?* FEE Stories. https://fee.org/articles/whos-to-blame-when-a-driverless-car-has-an-accident/#:~:text=The%20Parties%20to%20an%20Accident&text=Manufacturers%20could%20be%20liable%20in,inadequate%20service%20to%20the%20vehicle

Kerr, D. (2024, August 10). *California allows robo-taxis to expand and emergency responders aren't happy*. NPR KQED. https://www.npr.org/2023/08/10/1193272085/san-francisco-has-lots-of-self-driving-cars-theyre-driving-first-responders-nuts

Orlowski, A. (2024, January 29). *Self-driving cars are journeying into a moral black hole*. The Telegraph. https://www.telegraph.co.uk/business/2024/01/29/driverless-cars-ai-artificial-intelligence-tesla-apple/

Sabatier, Grant. “30 Gig Economy Statistics for 2024.” *Millennial Money*, 21 Mar. 2024, millennialmoney.com/gig-economy-statistics/#:~:text=A%20study%20from%20ADP%20Research,California%3A%2020%25

Sos, Z., Rasmus, A., & Torrez, A. (2023, October 2). *Cruise says it's not at fault for woman injured in San Francisco crash*. Fox KTVU. https://www.ktvu.com/news/cruise-says-its-not-at-fault-for-woman-injured-in-san-francisco-crash

---

Technical Exercises:

It may be tempting to believe that deep learning, or reinforcement learning, combined with coded rules, will be enough to train any vehicle how to drive on its own, and to do so more safely than human drivers given the response time of a computer versus a possibly distracted driver. However, as we’ve seen in the above case study and references, it can be difficult to decide the right thing to do in a given situation.

Pretend you are developing a rule-based algorithm for a driverless vehicle. In a Jupyter or Google Colab notebook,

- create a simplistic driverlessCar class that represents a vehicle that can be either manually operated or autonomous;
- define some attributes such as Make, Model, Year, Color, Speed, etc.. No need for this to be an exhaustive list, you can add to it later;
- define methods for enabling and disabling autonomous driving

Now, look at the following scenarios, and define one or more methods for the vehicle to react to each scenario. For each method, print to the screen if any human or animal is likely to be hurt or killed, and if any vehicles are likely to be damaged given the action that the vehicle is taking. After the driverlessCar class is defined, test each method below the class definition.

Scenarios:

1. A deer suddenly appears in front of the vehicle.

2. A deer suddenly appears in front of the vehicle. You are in the right hand lane and there is a vehicle directly next to you on the left side, and a vehicle behind you following very closely.
3. A deer suddenly appears in front of the vehicle. There are two small children in the vehicle with you. You are in the right hand lane and there is a vehicle directly next to you on the left side, and a vehicle behind you following very closely. It is raining heavily.
4. A deer suddenly appears in front of the vehicle. You will hit the deer no matter what you do.
5. A moose suddenly appears in front of the vehicle. There is a very large tree on the right hand side of the road, and a vehicle in the lane to the left of you. Nobody is following behind you very closely, but there is a very large heavy truck (not a pickup truck, a commercial truck) with a trailer following about 40 feet behind you going 30 miles per hour. You are going downhill at a 3% downgrade.

After implementing scenarios 1-5 above, comment on the complexity of engineering a driverless vehicle.