Synthesis of the Path-Planning of Autonomous Vehicles and Reinforcement Learning in University of Virginia Computer Science Curriculum. Analysis of the Failure of Uber's Autonomous Vehicles Program.

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Matthew Nicoud

December 6, 2023

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Benjamin Laugelli, Department of Engineering and Society

Briana Morrison, Department of Computer Science

Introduction

Autonomous vehicles are a topic of increasing importance and interest to society as companies such as Tesla work to develop cars that can be driven without any human intervention. The level of autonomy in a vehicle is determined using the Society of Automotive Engineers (SAE) classification system, which ranges from Level 0, No Driving Automation, to Level 5, Full Automation (Gordon, 2020). A vehicle is only considered truly autonomous when it reaches Level 3, Conditional Driving Automation, and very few companies have acquired regulatory approval to operate Level 3 vehicles on public roads. Even Tesla, perhaps the most well-known manufacturer of vehicles with autopilot, has been stuck on Level 2 autonomy for years, demonstrating the immense technical challenges of developing road-safe autonomous vehicles (Mulach, 2023). In order to prepare students to help advance automation in the motor-vehicle industry, the University of Virginia (UVA) has begun offering a Special Topics in Computer Science course titled Autonomous Vehicles: Perception, Planning & Control.

To further prepare students for the engineering projects that they will encounter in the workplace, I propose the synthesis of path-planning of autonomous vehicles and artificial intelligence reinforcement learning. As autonomous vehicles outside of education require the construction of a diverse network of technical, social, conceptual, and economic factors, it is necessary for students to understand the inner workings of successful network formation in relation to automotive projects and algorithms they are involved with. To examine the inner workings of the network of autonomous vehicles, I will draw on the science, technology, and society (STS) framework of Actor-Network Theory to analyze the failure of Uber's self-driving car during testing on a public road. Specifically, I will investigate how the relationship between technical and social factors involved with testing self-driving vehicles on a public road at high speeds contributed to the program's failure. These technical and social factors are directly correlated, involving the engineering design decisions behind Uber's self-driving algorithms, the regulations set by the local government, and the views of citizens in cities where self-driving cars are tested. Insights gathered from this STS project will further direct the exact manner of how reinforcement learning and path-planning of autonomous vehicles should be taught together at UVA.

Technical Project Proposal

With the introduction of the special topics course of Autonomous Vehicles at UVA, students are initially taught to write programs that use Light Detection and Ranging (LiDAR) sensors on 1/10 scale RC cars to have these cars act as reflex (reactionary) agents. These cars follow the wall and eventually make use of the "Follow The Gap" Algorithm to detect and avoid obstacles. This algorithm involves calculating a gap angle to go in the direction of the largest gap within its sensor range and is even used outside of education (Houshyari, 2021). Later in the coursework, students increase the autonomy of their vehicles by using LiDAR sensors and Vedder Electric Speed Controllers (VESC) to build their own occupancy grid map of a track, which is a grid where each cell is marked with a number that indicates the likelihood that the cell contains an object. Once the map has been established, a plugin is used to plan the path of the car without any further calculation by the student's program.

In the class Artificial Intelligence, students learn to solve non-deterministic search problems involving a grid called Markov decision processes (MDPs). A MDP is defined by a set of states (often in a grid), a set of actions determined by boundaries and obstacles, a transition function, a reward function based on the values in each state in the grid, and a start state (Kiran et al., 2020). To solve these problems, students are taught to use various reinforcement learning techniques such as value iteration and Q-learning to take the optimal path in the environment. The professor even uses navigating through a parking lot as an example of when reinforcement learning might be used to optimize the best parking spot to achieve, however there is no activity or assignment to apply the techniques to an autonomous vehicle. Instead, versions of these reinforcement learning techniques are implemented into assignments for students using an established Pacman game to visualize the paths taken. There is even a lecture on autonomous vehicles towards the end of the course, but there is no hands-on application of the techniques we learned throughout the class to provide a deeper understanding of the actual implementation.

Simply using a plugin to determine the path of an autonomous vehicle and copying pseudo-code into a few functions amidst a pre-established Pacman game package are insufficient ways of preparing students for the complex engineering challenges they will face in the workplace. Engineers that go on to work on autonomous vehicles that will be driving on public roads amongst civilians need to be experienced with how exactly the best route is determined, rather than relying completely on an outside resource that they have no in-depth knowledge of. By synthesizing reinforcement learning techniques with the path-planning of an autonomous vehicle, students will be able to test and perfect their understanding of how these techniques work in addition to learning how the path and speed of an autonomous vehicle is calculated. This not only includes taking the most efficient and optimal route, but also detecting and avoiding obstacles which is even more important to bring to projects beyond education. My proposed synthesis will directly translate to the autonomous vehicles industry where students might work, as motion planning directly makes use of various reinforcement and search algorithms such as Dijkstra's Algorithm, the A* Algorithm, and their derivatives (Gu, 2012). This will provide students who proceed to work with autonomous vehicles the specific skills they need to adjust path-planning and avoidance algorithms in the workplace to minimize regret, which in the real world can equate to loss of life when considering the speed of motor vehicles. I will make use of the syllabus, textbooks, and course materials from UVA's Artificial Intelligence and Autonomous Vehicles: Perception, Planning & Control courses to direct a complete and comprehensive course of action for how these two topics should be combined. To determine the exact way these topics can be combined in coursework, I will rely on Computer Science Curricula 2013 by ACM and IEEE, which explores the optimal ways to design university computer sciences courses.

STS Project Proposal

In 2015, Uber began working on self-driving cars through its Uber AV project. Soon after, officials in the Arizona government began removing regulations on autonomous vehicles in order to attract Uber and other competitors such as Waymo to test their self-driving cars in order to help the state's economy. Thus, Uber AV moved into Tempe, Arizona to run its tests in the dry conditions before elevating to areas with more consistent precipitation. Its tests involved collecting data from a set path and were conducted with an operator in the driver's seat who was meant to take over in case of emergency. In March of 2017, a vehicle collided with a car in Uber AV's fleet, where it was determined that the other vehicle was at fault for the accident. However, on the night of March 18, 2018, an Uber AV test vehicle struck and killed a pedestrian, Elaine Herzberg, who was pushing a bicycle across a four-lane road in Tempe (Wakabayashi, 2018). This marked the first known death to a pedestrian from an autonomous vehicle. After this event, Uber significantly cut back on its testing of self-driving cars on public roads, and eventually sold its autonomous vehicle business to a startup called Aurora Innovations in December of 2020 (Hawkins, 2020).

The discussion around the failure of Uber's self-driving fleet includes a variety of factors, including lack of public trust in the vehicles and operators, decreasing funds, increasing regulations, and legal disputes (Hawkins, 2019). All of these factors were involved with the end of Uber's autonomous vehicle business, but the common conception ignores the connection between these various actors that culminated in the demise of one of the frontrunners of the automation movement in the automotive industry. Placing the blame solely on individual aspects of Uber AV's history ignores the complex interaction between pedestrians, government officials, engineers, Uber drivers, Uber leadership, and test operators. For example, if the 2018 crash was the only reason for Uber to give up on its project, then it would have stopped investing its money into it immediately after. Instead, the company still projected that they could save money in the long run by replacing its drivers with self-driving cars and could even decrease the number of total collisions involved with its service (Jeffs, 2023). By focusing on a single event of a car failing to detect and avoid a person, we fail to gain a proper understanding of how one of the leaders in autonomous driving was entirely removed from the industry.

To frame my analysis of Uber AV, I will thus draw on the Actor-Network Theory (ANT) STS framework to offer an organized approach for investigating the relationships between all of these factors which altogether resulted in the beginning and eventual end of Uber's flagship program. ANT, developed by Michel Callon, Bruno Latour, John Law, and other STS writers in the 1980s, claims that all technical problems can be understood as a network of human and non-human actors. These networks are typically interpreted through network builders, from whose perspective we can analyze the constantly changing relationships between society and technology to view a network's success or failure (Cressman, 2009). To examine the interconnections between the network of human and non-human actors that resulted in the tragic accident in Tempe, I will draw on interviews from Arizona and Uber officials, direct accounts of the test operator Rafaela Vasquez during the March 18th crash, and legal proceedings from the death of Elaine Herzberg. I will also investigate litigation settlements reached between Uber and competitor Waymo which resulted in the sentencing of 18 months in prison to Uber engineer Anothony Levandowski.

Conclusion

My research demonstrates the relationship between the technical and STS projects proposed here to provide a plan that best prepares computer science students at UVA for the automotive industry. While the technical project proposes the synthesis of reinforcement learning techniques and path-planning of autonomous vehicles into a single piece of content, the STS project inquires into the potential consequences of testing imperfect automated technologies in the real world. Collectively, this brings forth a more in-depth way to address the sociotechnical challenge of teaching path-planning of autonomous vehicles with reinforcement learning and to further prepare UVA students for the projects and challenges they could face in the workplace.

References

Cressman D. (2009). *A brief overview of actor network theory: Punctualization, heterogeneous engineering & translation*. http://summit.sfu.ca/item/13593

Gordon, R. (2020, December 10). *Explained: Levels of autonomy in self-driving cars*. MIT CSAIL. https://www.csail.mit.edu/news/explained-levels-autonomy-self-driving-cars

Gu, T., & Dolan, J. M. (2012). *On-road motion planning for autonomous vehicles*. Springer. https://doi.org/10.1007/978-3-642-33503-7_57

Hawkins, A. J. (2019, November 19). *Uber is at fault for fatal self-driving crash, but it's not alone*. The Verge. https://www.theverge.com/2019/11/19/20972584/uber-fault-self-driving-crash-ntsb-probable-cause

Hawkins, A. J. (2020, December 7). *Uber's fraught and deadly pursuit of self-driving cars is over*. The Verge. https://www.theverge.com/2020/12/7/22158745/uber-selling-autonomous-vehicle-business-aurora-innovation

Houshyari, H., & Sezer, V. (2022). A new gap-based obstacle avoidance approach: Follow the obstacle circle method. *Robotica*, *40*(7), 2231-2254. doi:10.1017/S0263574721001624

Jeffs, J. (2023, September 12). *Autonomous vehicles will soon be safer than humans, some already are*. IDTechEx.

https://www.idtechex.com/en/research-article/autonomous-vehicles-will-soon-be-safer-than-humans-some-already-are/29840#:~:text=Another%20way%20to%20look%20at

Kiran, B. R., Sobh, I., Talpaert, V., Mannion, P., Sallab, A. A. A., Yogamani, S., & Pérez, P. (2020). *Deep reinforcement learning for autonomous driving: A survey*. ArXiv. https://arxiv.org/abs/2002.00444

Mulach, J. (2023, February 28). *Tesla admits its semi-autonomous driving tech is not the world's most advanced*. Drive. https://www.drive.com.au/news/tesla-full-self-driving-level-2-autonomous/

Smiley, L. (2022, March 8). *"I'm the operator": The aftermath of a self-driving tragedy*. Wired. https://www.wired.com/story/uber-self-driving-car-fatal-crash/

Wakabayashi, D. (2018, March 19). *Self-driving Uber car kills pedestrian in Arizona, where robots roam*. The New York Times. https://www.nytimes.com/2018/03/19/technology/uber-driverless-fatality.html