Maximizing Model Training with Minimal Resources
(Technical Paper)

The Integration of Social Factors in the Development and Implementation of Autonomous Vehicles
(STS Paper)

A Thesis Prospectus Submitted to the
Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia
In Partial Fulfillment of the Requirements of the Degree
Bachelor of Science, School of Engineering

Sarah Lei
Fall, 2019

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

Signature ____ Date ___________
Sarah Lei

Approved ____ Date ___________
James Cohoon, Department of Computer Science

Approved ____ Date ___________
Cody Fleming, Department of Systems and Information Engineering

Approved ____ Date ___________
Bryn Seabrook, Department of Engineering and Society

Introduction

Autonomous vehicles (AVs) were once envisioned to be technology of the future, but have now become a concrete reality. With autonomy software rapidly developing, it is easy to neglect to consider their implications on society. There have been critical casualties resulting from faulty logic in autonomy causing government policy to be enacted and enforced. AVs are changing how people think and react whether society is aware of it or not.

Stanford's Ethical Decision Making for Autonomous Vehicles states, "Human drivers navigate the roadways by balancing values such as safety, legality, and mobility. An autonomous vehicle driving on the same roadways as humans likely needs to navigate based on similar values" (Thornton, 2018). It is important as autonomy progresses to distinguish what those values are, how they are assessed, and how their implementation in systems are shaping society. Values are vastly impacted by culture and geographical location, but in what capacity are autonomous vehicles reflective of those values? The proposed STS paper aims to assess how social influences are being incorporated in the development and implementation of AVs. Furthermore, this paper will analyze those societal influences that have created a technologically deterministic future of AVs.

Inversely, autonomy is modeled after human behavior. The usability of data-driven decision-making solutions in safety-critical applications is limited by the absence of human-like inference and reasoning about decisions. Moreover, most models overlook the possibility of more than one feasible decision and end up optimizing on average behavior. The proposed technical project is expected to improve decision-making models by incorporating human expert knowledge and interpretability in the model by investigating how to maximize model training with minimal resources.

Technical Topic (Capstone)

Autonomous vehicles, also known as self-driving cars, "require methods that generalize to unpredictable situations and reason in a timely manner in order to reach human-level reliability" (Schwarting, Alonso-Mora, & Rus, 2018). These methods incorporate many elements in their logic that are developed through a deep learning based model. These models include end-to-end motion planning to learn a navigation policy in simulation from an expert operator. The knowledge gained from training is then transferred to real-world environments to perform target-oriented navigation and collision avoidance (Schwarting et al., 2018). The proposed technical topic focuses on the prediction of surrounding entities and using the information learned to form a decision. An explainable deep learning based model for predicting intent of maritime entities was developed as a part of Center for Visual and Decision Informatics (CVDI) Year 7 Project 'Improved Decision Making for Autonomous Systems' which models intent based on past observations, while also incorporating spatial interactions and temporal dependencies. The technical topic will support this ongoing project by providing a more efficient, optimized approach to training this deep learning based model.

The ongoing project aims to improve data-driven decision-making models for an autonomous system, particularly with the goal of making explainable predictions for intent of other surrounding systems without any direct communication. The goal of this project is to improve the usability of and trust in such data-driven decision-making models by integrating human experts' decisions and feedback into the model's knowledge pool. Overall, the project will create a model that jointly reasons about and predicts a distribution over plausible decisions and uses human expert knowledge to generate an optimized decision. The application will focus on autonomous vessels at sea, with the goal of broad applicability to other domains.

While the goal of this project is to construct an explainable decision-making process inclusive of human expert feedback for autonomous navigation in complex environments, the focus of this technical topic will be maximizing model training using minimal resources. The construction of the goal decision-making process relies on deep learning models that emulate human navigation. These models require optical data to assess the surrounding environment to make an informed decision. To acquire the optical data, perception systems integrate "many heterogeneous sensors with significant communications and computation bandwidth to capture and process high-resolution, high-rate sensor data" (Leonard et al., 2008).

However, current hardware techniques of imaging drive a high cost and place limitations on resolution enhancement (Yue et al., 2016). This expensiveness and limited enhancement leads into the motivation for super-resolution: the process of upscaling or improving the details within an image by often using low-resolution images as input upscaled to a higher-resolution (Thomas, 2019). Currently, specific applications of super-resolutions include surveillance, remote sensing imaging, and biometric information identification (Yue et al., 2016). In addition to providing a way to construct a higher quality image from one that never existed or has been lost, super-resolution paves way for better perception for AVs thus better decision-making. However, "more advanced, adaptive, and faster methods with extensive applicability are always desirable" (Yue et al., 2016).

This technical topic will focus on the applications of automated machine learning in super-resolution models. Automated machine learning implements a "survival of the fittest" mentality. This process makes real-world models more efficient by running systematic processes on raw data and selecting models that pull the most relevant information from the data.

Essentially, what automated machine learning does is “automatically choose[s] a good algorithm and feature preprocessing steps for a new dataset at hand…” (Feurer et al., 2015).

In the scope of automated machine learning, the technical topic will center around the use of Rivanna, the University of Virginia’s High-Performance Computing (HPC) System (Research, 2019). Used for computational research, this supercomputer has over 8,000 cores and 8PB of various storage. In addition, Rivanna offers parallel partitioning service. To perform maximal training with minimal resources, this service will be used in the parallelization of single-threaded TensorFlow super resolution models on a Central Processing Unit (CPU) to investigate the differences among those models and evaluate their performances. Instead of executing instructions one step at a time, parallelization allows the use of two or more processors (cores, computers) in combination to solve a single problem (Stout, n.d.). The automation of this process allows tracking of the best performing super resolution models thus yielding maximal training with minimal resources.

STS Topic

In 2018, Uber released an autonomous car that hit and killed a pedestrian due to a technological failure in recognition (Lekach, 2018a). Further investigation found that the driver had been streaming episodes on YouTube, Netflix, and Hulu leading up to the crash. This failure in recognition leads to the questioning of what algorithms were used to detect objects – namely, what did the algorithm lack or fail to consider that led to a critical casualty? The death of the pedestrian, Elaine Herzberg, was the “first attributed to a self-driving vehicle and prompted significant safety concerns about the nascent self-driving car industry, which is working to get vehicles into commercial use” (Shepardson, 2019). This event prompted the U.S. National Transportation Safety Board (NTSB) to investigate with the intent of issuing safety recommendations to prevent similar crashes (National Transportation Safety Board, n.d.). The crash was one of the first that instigated the incorporation of social influences into the development of AVs.

Similarly, in 2018, Tesla's Model X fatally crashed into a highway barrier while in Autopilot mode (Rosenberg, 2018). Moments before the crash, the driver's hands were not detected on the wheel even after the vehicle issued an audio "hands-on" warning and multiple dashboard warnings. According to Tesla, human error is certainly at fault (Rosenberg, 2018). In 2016 alone, the U.S. National Highway Traffic Safety Administration tracked a total of 37,461 lives were lost in crashes on U.S. roads that year. Driver error (of the human variety) is involved in 90 percent of accidents (National Highway Traffic Safety Administration, 2017). One of the greatest challenges in AVs is that they can never encounter all possible traffic situations, but cars can learn from their mistake and crashes (Lekach, 2018b)

In acknowledgement of the recent fatal AV crashes, development of AVs continue to integrate more features that handles task of driving users do not want to or cannot do for themselves. What started as simpler features, like cruise control and antilock breaks, have evolved into blind spot detection and advanced driver assistance components, such as self-park (National Highway Traffic Safety Administration, 2019). AVs have the potential to remove human error from the crash equation, which will help protect drivers and passengers, as well as bicyclists and pedestrians. Safety is just one of the social influences being integrated into autonomous systems.

> On the other hand, recent advancements toward partly/fully automated vehicles are poised to revolutionize the perception and utilization of travel time in cars, and are further blurring the role of travel as a crisp transition between location-based activities. (Malokin, 2019).

Society's desire to evolve driving from a mundane, repetitive task to an opportunity for productivity has changed how AVs are designed and what features are to be added. The progression has changed how society views travel time and driving in general. This view has evolved from its earlier motivations to lessen the amount of decisions a user has to make such as cruise control. Now, AV applications have expanded to eliminating driver fatigue the key factor in the 4,000 deaths caused by truck accidents each year (Protiviti, 2017). The U.S. Congress is already pushing for stricter trucking regulation in 2016 that will limit operating times for truckers (Lardner, 2019). Although stemming from a safety concern, this case introduces the element of ease. Societal demand has increased to reducing the labor force, and development of AVs have been molded to follow those demands by getting closer and closer to fully autonomous vehicles. These vehicles are labeled as Level 5, defined by The Society of Automotive Engineers (SAE) as the highest, which means the car is able to drive itself in any situation and condition – or is fully autonomous. Level 5 drastically differs from Level 0, meaning no vehicle automation with the driver fully in control (National Highway Traffic Safety Administration, 2019).

> University of Virginia professor Madhur Behl, a computer science professor who researches high-performance autonomous racing cars that are a tenth the size of a normal car, is confident in the cars' abilities to be agile and safe. "Overall, these cars will prove to be safer than human counterparts," Behl said. (Lekach, 2018).

Increasing societal demands concerning safety, ease, and efficiency are being met as AVs approach SAE's Level 5 categorization. Moreover, these present day societal demands have created a technologically deterministic future of autonomous vehicles. Once a decision-making footprint is set, it will hard to change and converge that logic to be ubiquitous among manufacturers of AVs – emulating technological momentum.

Autonomy is both directly and indirectly modeled after human behavior. How humans act in differing environments is one of the many behaviors autonomy attempts to emulate. Society and its values then play a large role in the decision-making process for autonomous vehicles. Elements of social determinism are showcased in the idea of using social norms and human emotions to improve the collision avoidance of AVs. For example, the standard agent modelling tool Netlogo is utilized to simulate the artificial society of AVs to show the benefits of modeling autonomy after humans (Riaz, 2018). As such, the STS Social Construction of Technology framework will be investigated and applied. Additionally, as the deterministic future of AVs will be researched, the STS Technological Determinism framework will also be used. There are several key players within the network of AVs, such as government, public, engineer, and animals. Due to the interplay of these factors in autonomous logic, Actor-Network Theory will also be considered while conducting the STS research. The desire of autonomous vehicles to provide simplicity and minimalism is translated into its appearance and design. For example, Tesla's Model X includes retractable door handles, a compartmentalized dash board, vehicle tracking, AutoSteer, and AutoPark (Ingle & Phute, 2016). Humans have evidently influenced certain technologies to look a certain way and to portray a certain lifestyle.

Research Question and Methods

The research question is: How have societal demands and influences been incorporated into the development and implementation of AVs? AVs have been rapidly progressing while be modeled after human behavior and designed for societal demands. It is crucial to assess how these demands affect the development of a state-of-the-art technology, and how that might in turn create a futuristic technological determinism of AVs as this technology has changed and might continue to change human behavior. These demands may include negative effects that would be useful to address now and exclude them from autonomy logic. In any case, society must analyze what demands have shaped AV implementation and how those demands translate over into technology that changes how society thinks and behaves.

This topic will be analyzed through documentary research methods, surveys, policy analysis, and historical case studies. Although AVs are a new technology, there are many articles that explain how AVs make decisions and what patterns they follow. Certain papers publish the policies AVs follow or should follow, so policy analysis will be used on these manuals to assess the current guides in place. This data collected will serve as a foundation to better understanding the existing frameworks. Additionally, there have been many cases of failure in AVs in the news. These historical case studies will be used to understand how policy has changed in response to these events. Surveys will be used to analyze statistically humans trust and ethics regarding AVs. Collection of data will be centered around the decision-making processes for AVs modeled after human behavior. The data was further be organized by a mixture of date and values. As AVs are a fairly new technology, policies and cases of failure will be considered within a timeline. The logic implemented in autonomous systems will be categorized by values, such as efficiency, safety, or aesthetic.

Conclusion

The goal of the technical topic is to construct explainable decision-making process, inclusive of human expert feedback, for autonomous navigation in complex environments through the means of utilizing Rivanna, UVA’s High-Performance Computing System, to analyze and perform maximal model training with minimal resources. This optimized model training will be investigated through the parallelization of single-threaded TensorFlow super resolution models on a CPU to investigate the differences among those models, evaluate their performances, and choose the best performing model. By incorporating human knowledge and interpretability, the model will be able to better emulate human behaviors. This will contribute to a more accurate and comprehensive model.

The STS topic will analyze the integration of social demands and influences on the development and implementation of AVs. These influences center around safety, ease, and efficiency. Furthermore, the STS topic will analyze how these influences have created a technologically deterministic future for AVs. This topic will be relevant to discern how AV technology is being developed in regards to societal demands and how the implementation of those demands is driving how AVs influence society.

References

Feurer, M., Klein, A., Eggensperger, K., Springenberg, J., Blum, M., & Hutter, F. (2015). Efficient and robust automated machine learning. In *Advances in neural information processing systems* (pp. 2962-2970).

Ingle, S., & Phute, M. (2016). Tesla autopilot: semi autonomous driving, an uptick for future autonomy. *International Research Journal of Engineering and Technology*, *3*(9).

Lardner, R. (2019, July 1). Trump administration to ease truckers' drive-time rules. Retrieved from https://www.pbs.org/newshour/nation/trump-administration-to-ease-truckers-drive-time-rules.

Leonard, J., How, J., Teller, S., Berger, M., Campbell, S., Fiore, G., … Williams, J. (2008). A perception-driven autonomous urban vehicle. *Journal of Field Robotics*, *25*(10), 727–774. doi: 10.1002/rob.20262

Lekach, S. (2018a, June 22). *Here's the fundamental flaw with self-driving cars: Humans.* Retrieved September 27, 2019, from https://mashable.com/article/uber-self-driving-fatal-crash-streaming-hulu/.

Lekach, S. (2018b, April 12). *Tesla's Autopilot fails haven't shaken my faith in self-driving cars. Here's why.* Retrieved from https://mashable.com/2018/04/12/self-driving-cars-safety-autonomous/.

Malokin, A., Circella, G., & Mokhtarian, P. L. (2019). How do activities conducted while commuting influence mode choice? Using revealed preference models to inform public transportation advantage and autonomous vehicle scenarios. Transportation Research Part A, 124, 82–114. https://doi.org/10.1016/j.tra.2018.12.015

National Highway Traffic Safety Administration. (2017, October 6). USDOT Releases 2016 Fatal Traffic Crash Data. Retrieved from https://www.nhtsa.gov/press-releases/usdot-releases-2016-fatal-traffic-crash-data.

National Highway Traffic Safety Administration. (2019, August 12). Automated Vehicles for Safety. Retrieved from https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety#topic-todays-tech.

National Transportation Safety Board. (n.d.). *Preliminary Report Highway: HWY18MH010*. *Preliminary Report Highway: HWY18MH010*. Retrieved from https://www.ntsb.gov/investigations/AccidentReports/Pages/HWY18MH010-prelim.aspx

Protiviti Inc. (2017). The Evolution of Autonomous Vehicles. Retrieved from https://www.protiviti.com/US-en/insights/evolution-autonomous-vehicles.

Research Computing. (2019, June 30). Retrieved from https://www.rc.virginia.edu/userinfo/rivanna/overview/.

Riaz, F., Jabbar, S., Sajid, M., Ahmad, M., Naseer, K., & Ali, N. (2018). A collision avoidance scheme for autonomous vehicles inspired by human social norms. Computers & Electrical Engineering, 69, 690–704. doi: 10.1016/j.compeleceng.2018.02.011

Rosenberg, A. (2018, March 31). Tesla Autopilot was turned on when the fatal crash in California occurred. Retrieved from https://mashable.com/2018/03/31/tesla-model-x-california-crash-autopilot/.

Schwarting, W., Alonso-Mora, J., & Rus, D. (2018). Planning and Decision-Making for Autonomous Vehicles. *Annual Review of Control, Robotics, and Autonomous Systems*, *1*, 187–210. doi: 10.1146/annurev-control-060117-105157

Shepardson, D. (2019, October 17). NTSB to determine probable cause of fatal Uber self-driving crash. Retrieved from https://www.reuters.com/article/us-uber-crash-safety/ntsb-to-determine-probable-cause-of-fatal-uber-self-driving-crash-idUSKBN1WW2IR.

Stout, Q. F. (n.d.). What is Parallel Computing? A Not Too Serious Explanation. Retrieved from https://web.eecs.umich.edu/~qstout/parallel.html.

Thomas, C. (2019, March 13). Deep learning based super resolution, without using a GAN. Retrieved from https://towardsdatascience.com/deep-learning-based-super-resolution-without-using-a-gan-11c9bb5b6cd5.

Thornton, S. M. (2018). *Autonomous Vehicle Motion Planning with Ethical Considerations*. (Doctoral dissertation). Retrieved from https://stacks.stanford.edu/file/druid:xf021wy6841/ThorntonSarah_thesis_final-augmented.pdf

Yue, L., Shen, H., Yuan, Q., Zhang, H., & Zhang, L. (2016). Image super-resolution: The techniques, applications, and future. *Signal Processing*, *128,* 389–408. doi: 10.1016/j.sigpro.2016.05.002