Cultural and Organizational Impediments to the Adoption of Self-Driving Vehicle Technology

(STS Topic)

Conducting Driving Simulator Experiments to Gauge Passenger Trust in Autonomous Driving Systems

(Technical Topic)

A Thesis Project 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

David Bennett

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 _____

Approved _____ Date _____

Kathryn A. Neeley, Associate Professor of STS, Department of Engineering and Society

Approved _____ Date 11/26/2019

Lu Feng, Assistant Professor, Department of Computer Science

Introduction:

Driving is an essential part of everyday life in any developed country. Whether as a driver, passenger, or pedestrian, it is nearly impossible for one to avoid regularly coming into contact with cars. Given its pervasiveness, it is fairly surprising how unsafe driving continues to be, with automobile accidents consistently ranking among the top causes of death for people across the globe. Even as the incremental introduction of safety features such as the seatbelt, rear view mirrors, airbags, and more recently blind spot monitoring have brought down driving fatalities considerably, the automobile remains the most dangerous common mode of transportation (other than the motorcycle), with automobile accidents causing over 40,000 deaths and 2,000,000 permanent injuries annually in the United States alone (Vardhman 2019).

The chief cause of the continued danger of traveling by car is the requirement for a human driver, as evidenced by the fact that 94% of all automobile-related fatalities were caused by human error or impairment (Maddox 2018). Autonomous vehicles offer a solution to this problem, removing the human element of driving entirely in favor of a series of specialized sensors and decision-making neural networks which, when perfected, will be able to outperform a human driver in all possible driving situations. While the potential of AV technology to dramatically improve the safety and efficiency of traveling by car is immense, there are several factors which will complicate and impede the development and eventual widespread adoption of self-driving vehicles. In the STS portion of this prospectus, I will be comparing and contrasting these various impediments in order to determine which of them is likely to be the most influential in precluding the advancement of AV technology toward widespread use.

A particularly compelling hindrance to consumer acceptance of AVs is the trust, or lack thereof, that passengers have in self-driving vehicles. To evaluate passenger response to autonomous vehicles, my technical research group be conducted human trials with a driving

simulator to determine what behaviors of AVs cause passengers to distrust those systems. In the technical section of this prospectus I will examine the experimental design of the experiments conducted and compare it with similar experiments in the field. Ultimately, the deliverables of my research will hopefully grant a wider understanding of the potential pitfalls facing researchers on the frontier of self-driving technology and provide insight into methods of mitigating those problems.

STS Topic: Cultural and Organizational Impediments to the Adoption of Self-Driving Vehicle Technology

The issues which impede driverless vehicle technology from becoming widespread can be broken down into three rough categories. Firstly, technical concerns such as the reliability of AV sensors, the difficulty in formalizing driving as an algorithmic task, and the requirement for increased mapping and communication infrastructure are all significant impediments to the advancement of the technology. Secondly, cultural factors including consumer trust, or the lack thereof, in driverless systems, societal attitudes about car ownership, and media response to fatal accidents will all play a large role. Finally, there are organizational issues such as the continuing legal ambiguity of driverless vehicles, regulatory restrictions, and the continuing cost of autonomous driving packages for the consumer. While the addressing technical aspects of driving automation are absolutely instrumental to the successful adoption of the technology, in this section I will be focusing primarily on some select challenges from the cultural and organizational categories to compare and contrast the likely influence of each, as I will be exploring the technical category more in the technical portion of this document.

Although some crashes will still occur once cars are made fully driverless and perfected, AVs have the potential to eliminate up to 90% of traffic accidents (Gao et al. 2014). Even so, the

accidents that do still occur once drivers are removed from the equation will likely involve complicated ethical decisions that will need to be resolved by an unfeeling machine. In a scenario where an AV is on a collision course with a pedestrian in the road, and safe avoidance of the pedestrian is impossible, should the vehicle choose to stay the course, killing the pedestrian but preserving the life of its passenger, or veer off the road, saving the pedestrian and killing its passenger? Does this calculus change when there is a crowd of pedestrians rather than just one? Answers to this deeply philosophical series of questions—commonly called the trolley problem—would be absolutely necessary to program into the system, even if such situations arise infrequently. In answering these questions, the firms designing the algorithms controlling AVs will need to establish a code of conduct which is at once consistent, publicly acceptable, and acceptable to the consumer. According to the results of a study by Bonnefon et al. (2016), these three requirements for such a code may be at odds with one another. In a series of surveys conducted by that research group, a vast majority of respondents said that it was morally right for an AV to sacrifice its driver to save 10 pedestrians, however when asked whether they would buy a vehicle which was programmed to do such a thing, only 19 percent responded in the affirmative. This points to the conclusion that the societally acceptable utilitarian solution likely to be created by the legal system will not be acceptable to the consumers of AVs, which could cause serious economic problems for self-driving technology as it moves toward widespread use.

An added complication to the moral dilemmas behind AVs is the media response to fatal accidents involving autonomous vehicles. While it is an unsettled question whether or not negative headlines generate more interest than positive ones, news organizations do have a vested interest in portraying extraordinary and exceptional events rather than the more subtle and complex everyday happenings and socioeconomic trends of the world (van der Meer, et al.

2018). Because crashes such as these would be extremely rare once the technical systems behind self-driving vehicles is finalized and perfected, and because it's unlikely that the technical aspects of self-driving vehicles will be understood by the general public, the added unfamiliarity would make these events perfect headline fodder for news organizations. In the aftermath of a highly-publicized fatal accident involving an Uber vehicle striking and killing a pedestrian in March 2018 in Tempe, AZ, the percentage of drivers who stated that they would be afraid in a fully autonomous vehicle increased from 63% to 71% over a four month period from December 2017 to April 2018 (Stewart et al. 2019). Highly-publicized AV accidents like this one will inevitably continue to occur even as the vehicles become more and more optimized to perform on the roads, and each headline will serve to further divide public opinion over the balance of costs and benefits that driverless vehicles provide.

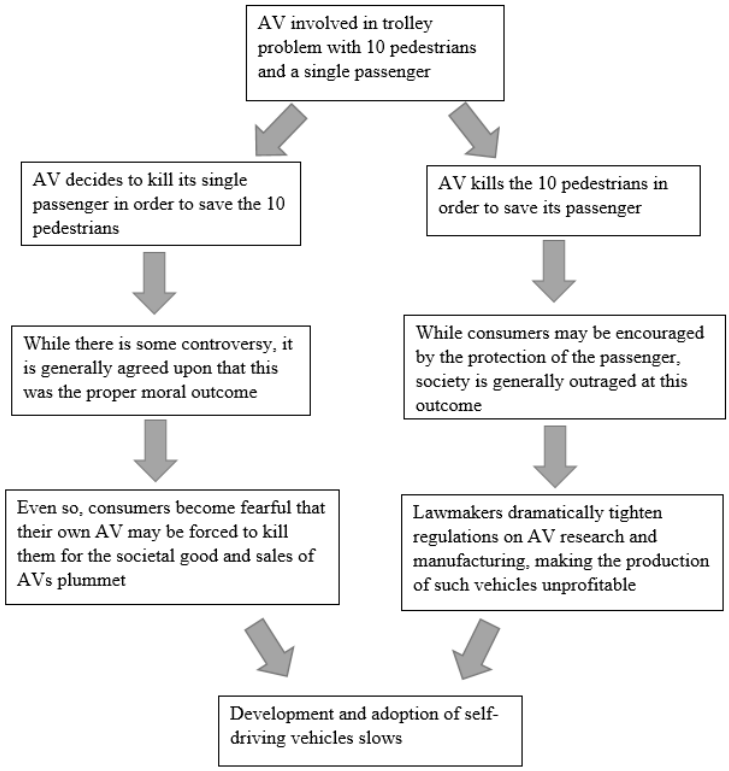

Figure 1: A flowchart outlining the possible outcomes of the trolley problem in practice

While it is difficult to pinpoint a single limiting factor which is the most challenging to overcome for AV research, the way forward is clear. Regardless of the direction in which future developments in the field go, transparency will be the most important requirement to ensure that the transition toward fully autonomous transportation is successful. More specifically, a public dialogue about the costs and benefits of the technology involving experts in the field, ethics experts, lawmakers, and consumers is paramount to ensuring that the pitfalls mentioned in this section are effectively avoided.

Technical Topic: Conducting Driving Simulator Experiments to Gauge Passenger Trust in Autonomous Driving Systems

Among the most prominent worries about autonomous vehicles is that the algorithms underlying them will not be able to safely handle the full range of contingencies that can occur when driving on roads filled with unpredictable human drivers. Indeed, if AVs are not able to accomplish this feat they will not only be unacceptable to the consumer but actively dangerous to their passengers. Even if an AV does perform safely in all possible situations, an influential factor in a passenger's level of trust in an autonomous system is the system's ability to reliably identify and alert the driver to incoming hazards on the road (this would be more important in an AV which allows driver intervention). Because of this, even a near perfectly-safe autonomous vehicle will be of no interest to consumers if it is unable to consistently alert the driver to hazards in a timely manner, regardless of the AVs ability to effectively handle the hazard itself. Alternatively, if passengers trust an autonomous driving system too much, they may not be prepared to reassume control of the vehicle in circumstances where human intervention is necessary.

In order to gauge the effects of various situations and factors on a passenger's trust in an autonomous driving system, we designed 16 test scenarios to be run on with human subjects in a Force Dynamics 401CR high-fidelity driving simulator. From an experimental design standpoint, using a driving simulator with physical feedback is a good middle ground between simply showing videos of scenarios to participants without feedback and using a fully-functional prototype autonomous vehicle. In the case of videos, an experiment without physical simulation to the movements of the vehicle would not do enough to immerse the participant in the experience, and thus would produce unreliable trust responses. On the other hand, using an autonomous system installed in a real vehicle would be both very expensive and would not allow our experiment to test a wide enough range of possible driving scenarios in the time scale of a single experiment.

In each test, the car would begin in autonomous mode, using a custom-built lane-keeping and self-driving algorithm, and would drive on a certain route through the course. Periodically, the simulated vehicle would encounter a randomized obstacle or hazard, such as an improperly crossing pedestrian or bicyclist in the middle of the lane. If the vehicle was still in autonomous mode, it would attempt to avoid the obstacle, and would either sound an alarm alerting the driver to the hazard or fail to alert, constituting a missed alarm. Throughout the experiment, the subject is able to continuously report his level of trust in the system on a scale of one to five where five was the highest confidence and one is the lowest confidence. For half of the tests, the participant had the option to switch the car into and out of manual driving mode at will. While driving in manual mode, the trust rating is automatically set to zero, and the driver was able to freely control the vehicle with the simulator's steering wheel and pedals.

Figure 2: Each participant was seated in the driving simulator and outfitted with a series of physiological monitoring devices such as an eye-tracker, galvanic skin response (GSR) sensor to track stress responses on the skin, and a photoplethysmography (PPG) sensor to track heart rate. Additionally, during each test the participant reported their trust electronically using the graphical user interface (GUI) shown above, while receiving information about the vehicle's status through that same GUI. Taken from the published paper from the research project (Sheng, et al. 2019)

Our research is distinct from the related inquiry conducted by Haeuslschmid et al. (2017) in that although their research also used human participants and investigated trust as the dependent variable, that study specifically tested the effect of different visualizations such as a digital chauffeur avatar or a world in miniature on passenger trust. Additionally, their experiment used a stationary Volkswagen Passat with driving scenarios projected onto a screen in front of the vehicle as a method of testing, meaning that physical feedback was not a factor in the experiment. Finally, that experiment gauged participant trust by polling each subject before and after the test, while our experiment allowed the participant to continuously report his trust in the system.

The main purpose of this method of inquiry was to directly expose a collection of subjects (potential passengers in future autonomous vehicles) to a series of success and failure

scenarios that could occur in a real AV, and to gauge which of those scenarios have the largest positive and negative effects on the user's trust in that system. With this information, we hope to analyze the areas which are the most important for autonomous driving systems to be robustly programmed in to encourage their adoption, as well as the extent to which partial failures (missed/false alarms) affect a passenger's confidence in a self-driving system. These conclusions could be potentially beneficial to the developers of the next generation of self-driving neural networks by raising awareness of problem areas for those systems which would make potential consumers hesitant to adopt the technology.

Conclusion:

While AVs have myriad possible benefits, the obstacles to their adoption are fairly considerable. From cultural and social attitudes to economic factors, there are several compelling elements which make the emergence of self-driving vehicle technology anything but a certainty. While consumer trust in AV technology may not be the single most impactful deterrent to self-driving cars becoming widespread, the results of my research group's experiments with passenger confidence have the potential to provide considerable insight into the scenarios which are most likely to dissuade otherwise willing consumers from buying autonomous vehicles. Having found that a series organizational and cultural factors including the inherent ethical ambiguities of AVs, media responses to driverless vehicle failures, and the public reaction to such failures are the factors most likely to impede the progress of AVs into the mainstream, firms on the cutting edge of research in this area should place a heavy emphasis on transparency about the costs and benefits of the technology. With the results of my technical and STS exploration, the hope is that further research and testing into AVs will proceed in a direction that maximizes the many positive qualities of the technology while minimizing the negatives.

Word Count: 2165

References:

Barabas, I, Todorut, A., Cordos, N., & Molea, A. (2017). Current challenges in autonomous driving. *IOP Conference Series: Materials Science and Engineering 252.* Retrieved from https://iopscience.iop.org/article/10.1088/1757-899X/252/1/012096/pdf

Bonnefon, J., Shariff, A. Rahwan, I. (2018, June 24). The social dilemma of autonomous vehicles. *Science, Volume 352, pages 1573-1576.* Retrieved from https://science.sciencemag.org/content/352/6293/1573

Boudette, N. E. (2019, July 17). Despite High Hopes, Self-Driving Cars Are 'Way in the Future'. *The New York Times.* Retrieved from https://www.nytimes.com/2019/07/17/business/self-driving-autonomous-cars.html.

Gao, P., Hensley, R., Zielke, A. (October 2014). A roadmap to the future for the auto industry. *McKinsey Quarterly*. Retrieved from www.mckinsey.com/industries/automotive-and-assembly/our-insights/a-road-map-to-the-future-for-the-auto-industry.

Gruel, W. & Stanford, J. M. (2016). Assessing the long-term effects of autonomous vehicles: a speculative approach. *Transportation Research Procedia*, *Volume 13,* pages 18-29. Retrieved from https://www.sciencedirect.com/science/article/pii/S2352146516300035.

Haeuslschmid, R., von Buelow, M., Pfleging, B., & Butz, A. (March 2017). Supporting trust in autonomous driving. *Proceedings of the 22nd International Conference on Intelligent User Interfaces.* Retrieved from https://www.researchgate.net/publication/314522858_Supporting_Trust_in_Autonomous_Driving

Litman, T. (2019, September 25). Autonomous vehicle implementation predictions: implications for transport planning. *Victoria Transport Policy Institute.* Retrieved from https://www.vtpi.org/avip.pdf

Maddox, T. (2018, February 1). *How autonomous vehicles could save over 350k lives in the US and millions worldwide.* Retrieved from https://www.zdnet.com/article/how-autonomous-vehicles-could-save-over-350k-lives-in-the-us-and-millions-worldwide/.\

Seif, H. & Hu, X. (2016, June). Autonomous driving in the iCity—HD maps as a key challenge of the automotive industry. *Engineering, Volume 2,* pages 159-162. Retrieved from https://www.sciencedirect.com/science/article/pii/S2095809916309432

Shaver, K. (2019, July 21). City planners eye self-driving vehicles to correct mistakes of the 20th-century auto. *The Washington Post.* Retrieved from

https://www.washingtonpost.com/transportation/2019/07/20/city-planners-eye-self-driving-vehicles-correct-mistakes-th-century-auto/.

Sheng, S., Pakdamanian, E., Han, K., Kim, B., Tiwari, P. Kim, I., & Feng, L. (2019, July 18). A case study of trust on autonomous driving. *arXiv*. arXiv:1904.11007. Retrieved from https://arxiv.org/abs/1904.11007

Stewart, L., Musa, M., & Croce, N. (2019, August 12). *Look no hands: self-driving vehicles' public trust problem.* Retrieved from https://www.weforum.org/agenda/2019/08/self-driving-vehicles-public-trust/

van der Meer, T., Kroon, A. C., Verhoeven, P., & Jonkman, J. (2018, January 16). Mediatization and the disproportionate attention to negative news. *Journalism Studies*, *Volume 20,* pages 783-803.

Vardhman, R. (2019, September 29). *Car accident statistics in the U.S. – 2019.* Retrieved September 30, 2019, from https://carsurance.net/blog/car-accident-statistics/.

Vincent, J. M. (2017, February 17). Top 10 Barriers to Self-Driving Cars. Retrieved September 30, 2019, from https://cars.usnews.com/cars-trucks/barriers-to-self-driving-cars.

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

World Health Organization. (2013). Global Health Observatory Data. Retrieved from https://www.who.int/gho/road_safety/mortality/en/