




Technologies for Safe & Efficient Transportation

THE NATIONAL USDOT UNIVERSITY
TRANSPORTATION CENTER FOR SAFETY

Carnegie Mellon University

UNIVERSITY of PENNSYLVANIA

Bystander Interactions with Failing Vehicle Autonomy

Jan. 1, 2017 – Aug. 31, 2018

FINAL RESEARCH REPORT

PI:

Aaron Steinfeld

<https://orcid.org/0000-0003-2274-0053>

Co-PI:

Jodi Forlizzi

Research Team:

Samadrita Das

Cecilia G. Morales

Selena Norman

Samantha Reig

Contract No. DTRT12GUTG11

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Problem

An important human factors challenge for autonomous vehicles (AVs) will be to find appropriate and acceptable methods for autonomous vehicles to express failure states to other drivers, bicyclists, and pedestrians. While efforts are underway to examine safe ways to hand control from the vehicle to the driver, work on alerting bystanders is sparse (see Reig et al, 2018 for details). It is important to address this challenge for two reasons. First, bystanders who are confused by the autonomous vehicle's actions are more likely to react incorrectly and therefore risk a collision. Second, vehicles that are viewed to act inappropriately will produce a negative impact on the vehicle brand and autonomous vehicle adoption as a whole. This project focused on the pedestrian-AV interaction challenge and explored forms of communication that are generalizable across autonomy type and brand.

Drivers currently express vehicle malfunction and emergencies in ways that are usually effective and safe due to driver training and road user experience. Related work on human interaction with autonomous mobile robots (e.g., Desai et al, 2013) suggests that feedback about robot system confidence is very important in producing timely and correct intervention by the novice users. Therefore, it is reasonable to believe that user friendly and rapidly understandable failure alerting methods exist for autonomous vehicles. However, it is important to also identify methods that are perceived to be appropriate.

There are cases in the automotive domain where warnings of failures are successfully viewed as appropriate and accepted by neighboring road users. For example, a driver who turns on their hazard lights and makes their way to a shoulder is generally perceived as being safe and given adequate room to maneuver. An autonomous system that uses similar methods to signal a desire to reach the shoulder is also likely to be tolerated. This is a logical action to take if the driver is slow to assume manual control. However, it is unclear how this will be perceived if the driver assumes manual control mid-action and resumes normal driving. Will this be viewed as safe? If not, is it possible to introduce other forms of information to neighboring road users to achieve the perception of safety? Should the driver be forced to wait until they reach the shoulder or right-hand lane before assuming manual control? Is there an equivalent to hazard lights that imply an autonomy state between normal operation and emergency?

Relevance to Stakeholders

Prior work suggests that an autonomous vehicle's brand will have a significant impact on the trust people place in the vehicle's autonomy (Carlson et al, 2014). This was associated with the perceived quality of the system and the competency of the engineering team. Therefore, it is important for autonomous vehicle companies to not just deliver safe vehicles with traditional assurances of high quality, but also maintain the perception that their vehicles behave appropriately and understandably. Behaviors and alerts that are perceived to be strange, unpredictable, and unexpected will have a negative impact on product success. Adoption is important from a societal perspective since the path to fewer crashes requires vehicle autonomy.

This concern is similar to prior work by the project lead on rear-mounted warnings to alert oncoming vehicles of dangerous rear-end approaches (Krishnan et al, 2001). The automotive manufacturer that sponsored the work was specifically concerned that a rear-facing warning that issued too many nuisance alarms would damage the company's brand. An interesting aspect of this project was that the sponsor did not want public dissemination of their involvement for various external relations reasons. We suspect similar concerns are present related to the work proposed here. Therefore, the team opted to not partner with specific organizations, unlike what is typically done in the T-SET Consortium, thus allowing partners to be insulated from direct connections to the project.

Approach

The team used methods that allow rapid exploration of suitable approaches. While on-road and laboratory testing is needed for conclusive results, it is important to narrow down possible approaches using cheaper, less time consuming methods. In this work, the team focused on small-scale interviews with regular road users to identify potential problems with the scenarios. These were used to both identify the core issues facing pedestrian-AV interaction (Field Study) and to assess preliminary design ideas (Design Study). For the latter, the team used traditional interaction design techniques to develop the information display types. The candidate displays were examined using small-scale interviews with regular pedestrians in naturalistic settings. The data for both efforts was analyzed with an eye towards steering the community on information displays and broader outreach needs by AV companies and policymakers.

This project ran in parallel with a grant from the National Science Foundation on human interaction with failing robots (IIS-1552256)¹. Similar findings were observed in both efforts, suggesting there may be generalizable lessons and recommendations for human interaction with autonomy within a physical form.

Field Study

This study is extensively described in Reig et al (2018), so we provide a short summary here. This work was done in collaboration with Jodi Forlizzi.

Methodology

We conducted an interview-based field study of AVs with pedestrians in nine settings around Pittsburgh, PA where autonomous Uber vehicles were routinely being tested at the time. We interviewed 32 participants over ten weeks across summer and winter seasons. All interviews occurred prior to the fatal crash in Arizona.

The interview questions were designed to gather insight on (1) whether pedestrians understand the operation and capabilities of AVs, (2) whether pedestrians trust AV technology, (3)

¹ http://www.nsf.gov/awardsearch/showAward?AWD_ID=1552256

pedestrian perceptions of AVs and artificial intelligence (AI) in general, and (4) whether brand impacts understanding, trust, or perception of AVs.

Interviewers waited near intersections where AVs were spotted regularly. When an AV and a pedestrian were both in the vicinity of an intersection and the pedestrian was not talking on a cell phone or seemingly in a rush, the interviewer(s) intercepted the pedestrian for an interview.

Findings

Data was collated, transcribed, and analyzed. Five core themes were drawn from the participant responses:

- Lack of awareness leads to mistrust and made-up explanations for the behaviors of AVs.
- Belief that AV technology is not yet mature leads to a lack of perceived trustworthiness of AVs.
- An “innocent until proven guilty” mentality is prevalent with respect to AVs and trust.
- People associate AVs with AI, and both are perceived by some to be detrimental to society.
- There are brand associations, both positive and negative, that affect the perceived trustworthiness of AVs.

These five themes have direct implications on AV deployments, with a clear need for better transparency by industry on how AVs work, what they are capable of, and what benefits they may offer individual users and society as a whole. It is also clear that favorable perceptions of company brand can produce favorable perceptions of the AVs.

This work was presented at a peer-reviewed conference. Findings from this work, and copies of paper (Reig et al, 2018), were also conveyed to multiple AV companies. In some cases, AV company personnel also had follow-up meetings with representatives of the team.

Display Study

Methodology

This work centered on an independent study by a Masters student (S. Das) under the guidance of Steinfeld and Forlizzi. This project focused on forms of communication between pedestrians and AVs. A set of designed interactions were prepared and tested in a natural setting in Pittsburgh.

Bystander Interactions with Failing Vehicle Autonomy

Information displays were developed using a programmable LED strip and a tablet device (iPad). Both were positioned at the bottom of the windshield of a conventional car that was parked on a busy pedestrian road. Figures 1 and 2 show two example displays. Figure 1 shows red LEDs and “Please Wait” as an attempt to tell pedestrians to not cross. A second state had a “Safe to Cross” with all LEDs lit green. The third state was identical to the prior one, except a subset of LEDs tracked the pedestrian as they crossed (Figure 2). These were inspired from meetings with AV stakeholders, concept cars, and other research in this arena. Four participants were interviewed for each of the three states (12 participants total).



Figure 1. Warning to not cross in front of the AV



Figure 2. Pedestrian crossing, LEDs indicate AV is tracking the pedestrian's motion

The team also examined forms of communication humans could express to AVs. We opted to explore hand gestures, based on earlier work by Reissner (2007). We specifically looked at hand gestures for Stop (flat hand towards car), Hailing a shared car/taxi (open palm towards car), and Slow Down (flat hand waving down). Ten participants were interviewed on a city street to explore understandability and generalizability.

Findings

This work revealed numerous areas of possible confusion in bystanders about the AV's intent and what they were supposed to do. For example, do red lights on the windshield indicate the vehicle will stay stopped, under autonomous mode, or that pedestrians should stop? One participant even associated the red LEDs as decoration. Only 3 of 4 participants correctly interpreted the moving LED segment, with the fourth perceiving it as a countdown of the time left for them to cross. These poor results suggest a clear safety risk during intersection crossing.

The gesture exploration results suggested that the three gestures were widely associated with the intended meanings, suggesting promise in pedestrian-to-AV communication. There was some concern (4/10) regarding needing to learn a new gesture language, which suggests adoption and utilization of current gestures is important.

One of the more interesting quotes from a participant reflected the challenge of command vs. information, "I don't want the car to tell me what to do. It should tell me what it is doing and I want to make my own decision." Similar to the Field Study, there was general confusion about the intent of the vehicle.

Since this was preliminary work, no paper was published on this study. However, these findings have been discussed with industry and other researchers attempting to run similar studies at a larger scale in more controlled conditions (MIT, VA Tech, etc). We expect lessons from this study to be reflected in their work.

Outcomes and Outputs

Research Summary and Technical Recommendations

The team aggressively disseminated this project's findings to industry and policy makers through small meetings, presentations at conferences, and a high quality conference publication. These insights will help inform research, development, and policy regarding how AVs and associated systems convey failure and status.

At the top level, the general public has many misconceptions about AV capabilities and how AVs perceive and act on the world. Greater transparency and outreach to the public by the AV industry is needed in order to reduce confusion and support pedestrian decision making. All of our data collection was in Pittsburgh well after the public deployment of the Uber AVs.

Therefore, participants were past the novelty window and there had been time for knowledge to disseminate through the news, word of mouth, and other forms. Some participants had even ridden in an Uber AV prior to their interview. Therefore, the general confusion about AVs cannot be attributed to limited exposure.

Likewise, more knowledge is needed on effective forms of communication between AVs and pedestrians. The larger scale research efforts underway (e.g., VA Tech) will hopefully shed light on possible solutions. Our initial findings suggest there may be value in displays that convey AV intent rather than command pedestrian action.

While the “innocent until proven guilty” perspective observed within pedestrians may now be shifted due to recent crashes, it may be possible to leverage the notion that AVs are still learning. This may help properly calibrate pedestrian trust in vehicle capabilities.

Building Capacity

Multiple students on this project have joined industry, thereby supporting additional knowledge transfer. The remaining student is pursuing a PhD and has incorporated aspects of this project into her long-term research agenda. All of the students who worked on this project, either funded or through class projects, were women:

- Samadrita Das, a Masters in Human-Computer Interaction student, participated in this project as an independent study. She was the primary student for the Display Study. She has since graduated and joined Google.
- Cecilia G. Morales, a Masters in Robotics student, participated in in the Field Study and has since graduated and joined iRobot.
- Selena Norman participated in in the Field Study while obtaining a Bachelor of Design with an Industrial Design concentration. She is now a freelance designer.
- Samantha Reig is a PhD student in Human-Computer Interaction. She participated in in both studies. She is continuing to conduct research on human interaction with autonomy.

Outreach & Dissemination

Steinfeld was also active in many efforts to translate knowledge learned from this project to consumers, students, decision makers, and industry. This included quotes in articles (e.g., Business Insider, LA Times, Popular Science, etc), participation in panels (e.g., ATI21 policy event), discussions with policymakers (e.g., Department of Labor, SAE, Government Accounting Office, etc), and tours for K-12 students.

References

Carlson, M. S., Desai, M., Drury, J. L., Kwak, H., & Yanco, H. A. (2014). Identifying Factors that Influence Trust in Automated Cars and Medical Diagnosis Systems. Proceedings of the AAAI Spring Symposium on the Intersection of Robust Intelligence and Trust in Autonomous Systems.

Bystander Interactions with Failing Vehicle Autonomy

Desai, M., Kaniarasu, P., Medvedev, M., Steinfeld, A., & Yanco, H. (2013). Impact of robot failures and feedback on real-time trust. In Proceedings of the 8th Annual ACM/IEEE International Conference on Human-Robot Interaction.

Krishnan, H., Gibb, S., Steinfeld, A., & Shladover, S. E. (2001). Rear-end collision-warning system: Design and evaluation via simulation, Transportation Research Record - Journal of the Transportation Research Board 1759, 52-60.

Reig, S., Norman, S., Morales, C. G., Das, S., Steinfeld, A., & Forlizzi, J. (2018). A Field Study of Pedestrians and Autonomous Vehicles. In Proceedings of the 10th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI '18). ACM, New York, NY, USA, 198-209. DOI: <https://doi.org/10.1145/3239060.3239064>

Reissner, U. (2007). Gestures and speech in cars. In proceedings of Joint Advanced Student School (JASS).