UC Berkeley

UC Berkeley Previously Published Works

Title

Towards trustworthy automation: User Interfaces that convey internal and external awareness

Permalink

<https://escholarship.org/uc/item/2rz9q652>

ISBN

9781509018895

Authors

Rezvani, T
Driggs-Campbell, K
Sadigh, D
et al.

Publication Date

2016-12-22

DOI

10.1109/ITSC.2016.7795627

Peer reviewed

Towards Trustworthy Automation: User Interfaces that Convey Internal and External Awareness

Tara Rezvani, Katherine Driggs-Campbell, Dorsa Sadigh,
S. Shankar Sastry, Sanjit A. Seshia, and Ruzena Bajcsy

Abstract—Suppose we are given an autonomous vehicle that has limitations, meaning that it may need to transfer control back to the human driver to guarantee safety in certain situations. This paper presents work on designing a user interface to assist this hand off by considering the effects of the expression of *internal* and *external* awareness. *Internal* awareness is the concept of knowing whether or not the system is confident in its ability to handle the current situation. *External* awareness is the concept of being able to identify the limitations as the car is driving in terms of situational anomalies. We conduct a user study to examine what information should be presented to the driver, as well as the effects of expressing these levels of awareness on the driver’s situational awareness and trust in the automation. We found that expressing uncertainty about the autonomous system (internal awareness) had an adverse effect on driver experience and performance. However, by effectively conveying the automation’s external awareness on the anomaly, improvements were found in the driver’s situational awareness, increased trust in the system, and performance after the transfer of control.

I. INTRODUCTION

With the development of Advanced Driver Assistance Systems (ADAS) and increasing levels of autonomy in driving, the role of the human driver is transitioning from truly “operating” the vehicle to monitoring the vehicle’s operation [1]. To many, this is an exciting paradigm shift, as automation has the potential to significantly impact safety on the road and how people utilize their time [2]. What was once commute time will now be time to relax and utilize in-vehicle entertainment, as was patented by Ford in early 2016 [3]. However, automation is not perfect. As it has been practiced with aircraft systems, a pilot is always required to be ready to take over control for difficult tasks, e.g. landing or take off [4]. Similarly, drivers are required to be ready to take over if the autonomous system detects difficult situations that humans would be better at handling [5], [6].

The major difficulty in these semiautonomous systems, where transfer of control is required, is the interaction and interface with the human driver, as there is often disparity

between how the system functions and how the human expects the system to perform. For instance, when adaptive cruise control was first in testing, there were discrepancies between what people perceived as safe, and what was truly safe [7]. In [8], the study found that some advanced safety systems are actually increasing collision rates. When the autonomy does not perform as expected, drivers tend to either abuse the functionality due to lack of understanding in communication or reject the system entirely due to loss of trust [9]. Further, there are a number of issues that can arise during transfer of control, (e.g. mode confusion or lack of situational awareness [10]). Many factors play into the success of the transfer of control including the driver’s response time [5], which is a function of the driver’s situational awareness; the environmental scenario that possibly caused the vehicle to hand off control; and the warning given before the human is handed full control. To avoid such mode confusions, and to tap into the potential benefits of ADAS and autonomous systems, the interaction and intercommunication between the autonomy and the human must be carefully understood.

Imagine an autonomous system that comes with certain limitations. For example, an autonomous vehicle may not exactly know how to properly navigate construction zones (e.g. the current Tesla Autopilot [11]), or an autonomous system may have limited functionalities (e.g. Volvo’s City Safety System [12]). Our key insight is the reasons for failure and uncertainty in autonomy arise either from complicated environment settings or difficulty in providing a safe and trustworthy autonomous controller. In this work, we propose a self-aware system that consists of *internal* and *external awareness*, meaning that it is able to detect these situations/anomalies and has a sense of confidence associated with these scenarios. Assuming the system is capable of identifying such occurrences, the question of our interest is: *How should the autonomous system effectively transfer control back to the human in a safe manner?*

In this paper, we present findings in developing a user interface (UI) that expresses the internal and external awareness of autonomy. While there are many methods for developing user interfaces, this work is motivated by Sturton et al., who present the use of formal methods to verify the correctness of UI’s for electronic voting machines [13]. To use these methods, first an expectation model is needed to identify what information the driver desires and expects of the system. Here, we present our first foray in designing a UI to gather such data, and give a starting point of a UI that aims to increase driver awareness and trust in the system.

This material is based upon work supported by the National Science Foundation Award ECCS-1239323 and by the Office of Naval Research MURI Award ONR-N000141310341 and DURIP Award N000141310679.

The study protocol was approved by the University of California Institutional Review Board for human protection and privacy, under Protocol ID 2013-07-5459. Each subject was first informed on the experimental procedure and written informed consent was obtained.

Tara Rezvani, Katherine Driggs-Campbell, Dorsa Sadigh, S. Shankar Sastry, Sanjit A. Seshia and Ruzena Bajcsy are with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA (e-mail: tara.rezvani@berkeley.edu, {krdc,dsadigh,sastry,ssehsia,bajcsy}@eecs.berkeley.edu).

This pilot study evaluates the effects of presenting information about the internal and external awareness of autonomy under various for transition conditions with a distracted driver. This evaluation includes analysis of user feedback via surveys; driver monitoring to assess situational awareness; and assessment of driver performance after assuming control.

Through these scenarios and user interfaces, this paper aims to identify *what* information should be included, and *how* it should be presented to a driver to enable the safest transition. The contributions of this work are as follows: (1) designing UIs that express internal and external awareness of the vehicle to assist transfer of control to humans under various anomalies; (2) assessing the effects of expressing internal and external awareness from these UIs; and (3) analyzing the effect of information on the driver's perceived and measured awareness as well as performance on the road.

The paper is organized as follows. First, an overview of related works is presented. Section III presents the methodology behind our UI design, in addition to the experimental setup. The findings of this study are presented in Section IV and discussed in Section V.

II. RELATED WORKS

In the work presented here, we evaluate UI design on distracted drivers. There is a rich body of research in human factors and human-computer interaction associated with how to best increase situational awareness in operators [14], [15].

An important factor in this study was “when” to send the official transition warning to the driver. Research has advocated for a 5 to 8 seconds period for a safe takeover interval in order to avoid obstacles of varying complexities [16]. Previous studies suggest that shorter takeover requests cause faster reaction times yet poorer performance in taking over control [17]. Advanced warnings before a takeover request were found to correlate with a high rate of successfully avoided collision [18], which is adopted in this work.

In addition to timing for warnings, audio and visual cues have been well studied. In [16], audio warnings were found to be sufficient to warn the driver of a transfer of control. De Waard et al. tested visual cues through a study that would notify drivers of a transition, which was unsuccessful in catching the driver's attention [19]. Noting the success of audio and visual warnings, the UIs to be presented use a combination of a bold visual warning and an audio cue.

Given that drivers will be more prone to inattention, many studies have investigated how different distractions and obstacles will affect the driver's ability to assume control of the vehicle [20]. The main measure was the driver's success in takeover versus the distraction they were engaged with at the time of the take over. Maltz et al. [21] studied how an unreliable warning system for transfer of control demonstrates the driver's dependence on the warning system to take over the vehicle. The study found that without the warnings, the driver was unable to safely transfer control from autonomous to human driving. Research by Young et al. suggests that the takeover is affected by limited cognitive resources that can get consumed by multitasking [22].

Walker et al. proposes that when drivers are given information about their surroundings, they spend less time scanning the environment and have more successful takeovers [23]. Therefore, we assess the type of information about the reason of transition that is needed to decrease this searching time.

In the field of human-computer interaction, many studies have found that it is important that the human has a shared mental model to easily gain insight on how the robot is faring [24]. Similarly, Takayma et al. has found that expressing awareness in terms of uncertainty and reactions to failures can improve people's perception of a robot, and increase the robot's readability [25]. In his study, Steinfeld has interviewed experts in the field of autonomous or semiautonomous mobile robot interface experience and concluded that users should be able to quickly and easily identify the health and status of the robot [26]. This motivated our examination of internal awareness as an expression of robot's self-confidence in a given situation.

This paper is motivated by and builds off of many of these related works. The primary focus of this study is to evaluate the automation's expression of awareness to a distracted driver, in terms of type of awareness (internal and/or external) and the level of detail presented (general or detailed warnings). Our goal is to assess how this information affects the driver's awareness and trust in the autonomous system, under different transitions conditions.

III. METHODOLOGY

As previously described, this study targets scenarios in which there exists a semiautonomous system that has sufficient functionality such that full attention is not required from the human (i.e. the driver is free to watch in-vehicle entertainment), but might have to transfer control back to the driver if its limitations are met. We assess different causes for transferring control and UI designs that present information about the autonomy's internal and/or external awareness. The goal is to better understand the optimal level of information required for ideal transfer of control.

The following subsections describe the UI design methodology, the parameters of the user study, and the experimental setup to test the UI in an autonomous framework.

A. User Study

The user study was designed to test users' reactions to multiple UIs paired with a various driving scenarios. Varying the scenarios gives a more wholesome picture of the UIs' effect on the overall performance. The driving scenarios were selected such that they were variable and comprehensive. The three general categories of scenarios that could cause the autonomous vehicle to transfer control to the user are described in detail, followed by the study parameters:

(1) *Baseline Transition*: This transfer of control occurs when the AutoPilot needs to pass control back to the driver even when no immediate danger is detected. This event might occur when an autonomous system has reached the end of a known route or area, triggering a transfer of control back to the human driver.

(2) *Static Anomaly Detected*: These transitions occur when an unexpected static obstacle is detected and the autonomous system is not confident in how to proceed. Examples of these events include unexpected construction or a trailer blocking most of the road.

(3) *Dynamic Anomaly Detected*: Similar to the previous scenarios, this transfer of control occurs when an anomaly is detected. However, the anomaly in this case is in the form of a moving obstacle. This could be a nearby car that is swerving or behaving in a peculiar manner, meaning that the AutoPilot cannot predict what the vehicle will do.

Manipulated Factors: We manipulate two factors: (i) *Internal Awareness*: The current confidence level of the vehicle in its actions during AutoPilot mode. We use an emoticon shown in the first column of Figure 1 to describe the confidence in actions of the autonomous car. The emoticon takes three different states, each representing a high, medium and low level of confidence in the current actions the autonomous vehicle takes. (ii) *External Awareness*: The state of the environment, i.e., how complex the current environment conditions are during the autonomous mode. For the state and complexity of the environment, we use the measures shown in the second and third row of Figure 1 to represent either a static or dynamic anomaly. The specificity of the anomaly is general between the static and dynamic anomalies in the second column, i.e., all anomalies are displayed using a similar icon, whereas in the third column each anomaly is separately designated.

Dependent Measures: For each factor above, and the combination as in the last column of Figure 1, we measure different values to study the effectiveness of each factor. First, we evaluate measures for driver performance for a short period of time after taking control. We consider the average of the human's throttle and braking input, as well as the difference between the human driver's trajectory and a nominal trajectory in the same short horizon after human

takes over control. Second, we consider the driver's gaze, i.e., how often the driver looks at the front screen or the UI, to estimate situational awareness. We also consider measures regarding safety and trust based on a user survey.

Hypothesis: Our hypothesis is that using UIs representative of the two manipulated factors will result in better performance demonstrated by the described dependent measures. The expression of both internal and external awareness will lead to better understanding of the autonomous system, which will improve trust, awareness, and performance. However, it is suspected that the complete UI will show lower metrics than the other UIs due to clutter.

Subject Allocation: We recruited 10 participants (4 female, 6 male) in the age range of 18-61. All the participants owned a driver's license. We used a within-subjects design, and counterbalanced the order of conditions.

B. User Interface Design

The experiment was structured such that each participant would experience five different user interfaces and three different control transfer scenarios. Each participant would experience the following user interfaces: Baseline, Internal Confidence, External - General Warning, External - Detailed Warning and Complete. The features included in each of the user interfaces are as follows:

Baseline (UI 0): This user interface displays the car's speed and the current controller of the vehicle (AutoPilot or Human). During the transfer of control, there is a 5 seconds and 1 second audio warning beep, each with a different pitch. Along with the beeps there is a visual timer indicating 5 seconds and 1 second remains until the switch, as shown in Figure 2. The speed, controller and timer warnings are standard across all of the following UIs. The red box in Figure 2 points to where new features in the following UIs are placed relative to the baseline features.

Internal Confidence (UI 1): This user interface includes all the features in Baseline. There is also a visualization of the car on the road and an indicator of the AutoPilot's confidence in its own actions as shown in the first column of Figure 1. A smile emoticon indicates that the car is confident in its maneuvers. A neutral emoticon indicates that the car is not in immediate danger; however, it is unsure of some future actions. A frown emoticon indicates that the car is unable to handle a situation and needs to switch control to the human.

External - General Warning (UI 2): This user interface shown in the second column of Figure 1 builds off of Baseline. Added to it is a visualization of the car on the

	UI 1: Internal Confidence	UI 2: External – General Warning	UI 3: External – Detailed Warning	UI 4: Complete
Baseline				
Static Anomaly				
Dynamic Anomaly				

Fig. 1: UI visualization for anomaly detection. Each row presents the different reasons of switching back control. Each column corresponds to the UIs visualization for each anomaly.

Fig. 2: Skeleton of the UI design. The red box can be modified depending on specific UIs used in the experiment.

Fig. 3: Experimental setup with the car simulator. The steering wheel, MS Kinect for driver monitoring, the UI, and the entertainment tablet are specified above.

road with a generic hazard warning in the general region where an anomaly is occurring.

External - Detailed Warning (UI 3): This user interface includes all the features in Baseline. There is also a visualization of specific warnings of imminent potential hazards, as shown in column three of Figure 1. Example warnings include the location of construction or an unpredictable vehicle on the road.

Complete (UI 4): This combines features from Internal Confidence and External - Detailed Warning. This includes everything in Basic, the emoticon health indicators from Internal Confidence, and the warning icons from External - Detailed Warning, as depicted in the last column of Figure 1.

C. Experimental Setup

In order to collect data, we used a human-in-the-loop testbed as shown in Figure 3 [27]. This testbed consists of a car simulator to recreate driving experiences for the driver, as well as sensors to monitor the driver. Data was collected using a Force Dynamics CR401 vehicle simulator, which provides a 120° visual display from the driver's perspective¹. This system has been integrated with PreScan software [28], which provides vehicle dynamics and customizable driving environments. The system ran the simulation and data collection at 60 Hz. The driver was monitored using a Microsoft Kinect 2.0², which was used to monitor the attention of the driver. The driver was monitored at a rate of 30 Hz, which was synchronized with the simulation data.

The user interface, displaying key information to aid the computer to human control transition, is displayed over the simulated dashboard of the vehicle. To distract/entertain the driver while being driving by the autonomous system, a tablet playing videos is attached to the frame of the simulator. The complete experimental setup is shown in Figure 3. Each of the five user interfaces is implemented in Python 2.7 using the software package TkInter as the main GUI library [29]. The complete system was implemented on a 2.3 GHz Intel Core i7 processor with 16 GB RAM.

One test course was created and used for all experiments. To ensure that the driver did not become too familiar with

the course, variations on the visual world were created (e.g., different landscaping). The transition times and locations changed in each test, so the driver could not anticipate when to take control. The autopilot was controlled via a basic path following controller that would attempt to maintain a speed of 15 m/s and stay in the center of the lane.

For each user interface, the three scenarios were tested twice in a random order, for a total of six trials per user interface. In each trial, the AutoPilot would be in control for anywhere from three to five minutes, during which the driver was asked to watch videos on the entertainment tablet³. Then, one of the three scenarios would trigger a transition and the UI would warn the driver five seconds prior to the driver taking control. For consistency between the anomaly scenarios, the hand off was timed to always occur when the driver has two seconds of time headway to the obstacle [30].

IV. RESULTS

Our study's results indicate that the External-Detailed Warning user interface was the most successful, followed by External-General Warning, Complete, Internal Confidence and Baseline UIs, respectively. An overarching result gathered from the data is that the vehicle's external awareness caused the driver to trust the vehicle more and thus had better performance during the transfer of control. This disproves the hypothesis that conveying internal awareness increases trust in this human-robot system. However, this supports the idea that effectively conveying external awareness will increase the driver's perceived awareness and trust in the system.

The following subsections present the findings of the user study, looking at the data collected from the user feedback from the surveys, driver monitoring, and trajectory information after the transition occurred.

A. Subject Feedback Survey

In order to gauge the user's level of trust and security towards the autonomous vehicles, we requested subjects complete quantifiable surveys. These surveys aimed to target the user's fluctuating feelings of trust and safety throughout the entire experiment. The subjects were surveyed before experiencing the simulations, between each user interface session and at the end of the entire experiment. The initial survey and final survey asked on 5-point scale (with the options ranging from strongly disagree to strongly agree) the following three questions:

- Do you feel safe in a human-driven car?
- Do you feel safe in an autopilot?
- Is an autopilot safer than a human?

The initial and final responses are illustrated in Figure 4. On the chart, -2 represents strongly disagree and 2 represents strongly agree. Our results show that subjects felt that autonomous vehicles were safer than human drivers after the final survey. Further, our results reveal that at the end

¹<http://www.forcedynamics.com>

²<https://developer.microsoft.com/kinect>

³Simple distraction tests were performed prior to the experiment to verify that three minutes of watching videos was enough time for the driver to lose interest in the autonomous vehicle. More details on the driver's level of engagement are provided in Section IV-B

of the experiment their experiences with our user interfaces influenced their feelings about the safety of autonomous vehicles. There was an overall increase in feelings of safety of semi-autonomous vehicles and a slight decrease in their thoughts about the safety of human-drivers.

Fig. 4: This graph compares initial and final ratings of the survey questions concerning safety of human driven and autonomously driven vehicles. The values were collected based on a 5-point ranking system, where negative values demonstrate disagreement, 0 is neutral and positive values demonstrate agreement.

In both the initial and final surveys, we asked subjects how important audio and visual warnings were to them to help aid the takeover transition. On average, the importance of audio warnings increased by 0.6 points, while visual warnings only increased 0.2 points. This indicates that the audio cues were found to be more crucial than initially expected.

In between each user interface run, the survey asked, on a 5-point scale, if they trusted the semi-autonomous vehicle and if they felt the semi-autonomous vehicle was aware of its surroundings. The result of averaging these scores across all user interfaces is plotted in Figure 5. There is a high correlation coefficient between the two measures ($R^2 = 0.967$), which indicates that subjects trust the semi-autonomous vehicle most when they felt that the vehicle was most aware of its surroundings. Fluctuating values of trust and awareness are a result of the different user interfaces displaying different amounts of information in various ways.

There is also an interesting trend across the UIs in how they affected the subject's awareness of the environment and feelings of trust, as shown in Figure 6. The graph demonstrates the feeling of awareness the drivers received from the interface and their feeling of trust in the vehicle, for each user interface trial (in sequential order). The figure demonstrates that subjects are in stronger agreement in UI 3 and UI 4 interfaces, as majority of responses indicated high trust and increased awareness. This demonstrates that as more and more information was given to the driver through each UI (in order from left to right) the driver gains more awareness and trust.

B. Driver Monitoring

Using the data collected from driver monitoring, two key types of information were identified: (1) subject distraction and (2) search time. Subject distraction was evaluated to

Fig. 5: This graph shows the high correlation between driver's trust of the autopilot and the driver's feeling that the autopilot is aware of its surroundings as demonstrated through the UI.

determine whether or not the driver was sufficiently invested in the entertainment provided. Using the head pose relative to the Kinect at each time instance, the subject was considered distracted if they were looking in the direction of the tablet. For each trial, the ratio of time spent engaged in watching the entertainment was computed and compared in Figure 7.

We note that in general, the drivers were sufficiently distracted with the entertainment, but also that they were most engaged with UIs 2, 3, and 4. We attribute this to the fact that subjects trusted these UIs more than the first UIs.

Search time is the amount of time it takes for a driver to check the UI and identify the reason for transition in the real world. This was calculated using the head pose and eye movement, by finding the time it takes the driver to look ahead and for eye movement to settle after the warning signal. In essence, this attribute is highly related to the situational awareness that the UI is providing the driver when needed. A low search time would imply that the UI was able to quickly inform the driver the reason for transition and making it easier for her to identify, while a high search time would imply that the driver spent most of her transition warning period searching the screens for the problem. This is visualized in Figure 8.

It can be observed that the minimum search time corresponds to UI 3, closely followed by UI 2 and 4. Considering the amount of information presented by each UI, it appears that UI 3 provides the optimal amount of information to increase the driver's situational awareness. Since UI 3 and 4 provide similar information, it's intuitive that these UIs give similar results. We hypothesize that the increase in search time was due to UI 4 being more cluttered than UI 3.

Although UI 2 provides little information, the subject feedback indicated that the large visual warning caught their attention with peripheral vision, before the audio warning signal sounded. This implies that the drivers were aware of an anomaly with slightly more lead time than the other UIs.

Fig. 6: The bubble plot shows the all subject responses to survey questions, where the size of the point indicates the number of identical responses. Each column corresponds to UI 0 through 4. The graph displays the Trust Score on the x-axis versus the agreement that the UI increased awareness on the y-axis.

Fig. 7: This plot shows the level of engagement in the entertainment throughout each test for each UI, quantified as the ratio of time spent watching videos. Pink lines show the engagement for each trial conducted. The solid black line shows the average engagement, and one standard deviation is shown in the dashed line.

Fig. 8: This plot illustrates the average time spent searching during the five second warning period before the transition of control, for each user interface. The average search time for each subject is shown in pink. The average for all subjects is shown in solid black, and one standard deviation is shown in dashed line.

C. Driver Performance

The subject's control inputs and resulting trajectories were collected to assess the quality of taking over control. A safe takeover is one that doesn't involve significant deviation from a nominal trajectory with minimal braking or acceleration. A nominal trajectory is one driven by an expert driver without any distractions or transfer of control.

Figure 9 shows the average brake and acceleration inputs of all users in the case of dynamic anomalies. Since humans tend to perform evasive maneuvers in these scenarios, these measures are representative of these complicated takeovers. As shown, UI 2 and 3 perform well as their average braking

Fig. 9: Distribution of driver input for various UIs in the case of dynamic anomalies, with mean indicated by a point. Fig. (a) shows the average throttle input for all users and UIs. Fig. (b) show the average brake input for all users and UIs.

and throttle is lower than the other UIs, as does UI 4, despite being cluttered. However, in the case of dynamic obstacles, UI 1 is not sufficient for a safe transfer of control as the values are significantly higher than the other interfaces.

Also we consider the difference between trajectory performed by the users and a nominal trajectory shown in Figure 10. As illustrated, UI 3 has the lowest of deviation from the nominal trajectory for the baseline and static anomaly case, where all the other UIs perform worse. UI 3 also performs reasonably well for the dynamic anomalies case. However, UI 2 and 4 perform similarly in this dynamic case. We believe the deviation metric is not as descriptive as the input metrics for dynamic anomalies as the human might need to perform complicated trajectories in such complicated scenarios to stay safe. From our experimental results, we conclude that *UI 3 is the most effective for transfer of control*.

V. DISCUSSION

In this work, we designed UIs to convey internal and external awareness in order to better engage the driver and safely transfer control. By conveying external awareness precisely, the driver's perceived and measured awareness increased, along with their trust in the system. Improved performance after the handoff was also observed. Contrary to our initial belief, conveying internal awareness made people feel uneasy, likely due to the safety critical nature of driving.

Fig. 10: Boxplot of difference in drivers' trajectory and nominal trajectory for all users, and all UIs in the case of baseline, static anomalies and dynamic anomalies. Mean is indicated by a point.

In future work, we aim to quantify performance by applying model checking techniques, verifying logical properties on performance models and comparing the probability of failure, similar to work in [31]. Our hypothesis is that the quantitative values will prove that the driver's performance increases when using succinct, formally proven UIs in addition to a positive qualitative response from surveys. In addition, we hope to explore different UI mediums, and an expanded user study in more realistic settings.

ACKNOWLEDGEMENTS

We would like to thank Vijay Govindarajan for his help with the driver monitoring system used in the driving testbed.

REFERENCES

- [1] C. P. Janssen and J. L. Kenemans, "Multitasking in autonomous vehicles: Ready to go?" <http://www.auto-ui.org/15/p/workshops/4>, 2015.
- [2] M. Bertoncello and D. Wee, "Ten ways autonomous driving could redefine the automotive world," 2015.
- [3] M. Cuddihy and M. Rao, "Autonomous vehicle entertainment system," Mar. 1 2016, uS Patent 9,272,708. [Online]. Available: <https://www.google.com/patents/US9272708>
- [4] D. A. Norman, "The 'problem' with automation: inappropriate feedback and interaction, not 'over-automation'," *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, vol. 327, no. 1241, pp. 585–593, 1990.
- [5] W. Li, D. Sadigh, S. S. Sastry, and S. A. Seshia, "Synthesis for human-in-the-loop control systems," in *Proceedings of the 20th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, April 2014, pp. 470–484.
- [6] S. Ghosh, D. Sadigh, P. Nuzzo, V. Raman, A. Donzé, A. L. Sangiovanni-Vincentelli, S. S. Sastry, and S. A. Seshia, "Diagnosis and repair for synthesis from signal temporal logic specifications," in *Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control*. ACM, 2016, pp. 31–40.
- [7] M. Goodrich, E. Boer, and H. Inoue, "A model of human brake initiation behavior with implications for acc design," in *Intelligent Transportation Systems, 1999. Proceedings. 1999 IEEE/IEEE/JSAT International Conference on*, 1999, pp. 86–91.
- [8] S. Karush, "Insurance claims data show which new technologies are preventing crashes," *Status Report*, vol. 47, no. 5, 2012.
- [9] L. Hawkins. (2010, March) The skinny on self-parking. [Online]. Available: <http://online.wsj.com/news/articles/SB10001424052748703734504575125883649914708>
- [10] J. Rushby, "Using model checking to help discover mode confusions and other automation surprises," *Reliability Engineering & System Safety*, vol. 75, no. 2, pp. 167–177, 2002.
- [11] M. Ramsey. (2015, October) Driver videos push tesla's autopilot to its limits. [Online]. Available: <http://www.wsj.com/articles/driver-videos-push-teslas-autopilot-to-its-limits-1445765404>
- [12] E. Coelingh, L. Jakobsson, H. Lind, and M. Lindman, "Collision warning with auto brake - a real-life safety perspective," *National Highway Traffic Safety Administration*, April 2007.

- [13] C. Sturton, S. Jha, S. A. Seshia, and D. Wagner, "On voting machine design for verification and testability," in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, Nov. 2009, pp. 463–476.
- [14] M. Cellario, "Human-centered intelligent vehicles: Toward multimodal interface integration," *IEEE intelligent systems*, no. 4, pp. 78–81, 2001.
- [15] A. Schmidt, A. K. Dey, A. L. Kun, and W. Spiessl, "Automotive user interfaces: human computer interaction in the car," in *CHI'10 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2010, pp. 3177–3180.
- [16] B. Mok, M. Johns, K. J. Lee, D. Miller, D. Sirkin, P. Ive, and W. Ju, "Emergency, automation off: Unstructured transition timing for distracted drivers of automated vehicles," in *Intelligent Transportation Systems (ITSC), 2015 IEEE 18th International Conference on*. IEEE, 2015, pp. 2458–2464.
- [17] C. Gold, D. Damböck, L. Lorenz, and K. Bengler, "take over! how long does it take to get the driver back into the loop?" in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 57, no. 1. SAGE Publications, 2013, pp. 1938–1942.
- [18] A. P. van den Beukel and M. C. van der Voort, "The influence of time-criticality on situation awareness when retrieving human control after automated driving," in *Intelligent Transportation Systems (ITSC), 2013 16th International IEEE Conference on*. IEEE, 2013, pp. 2000–2005.
- [19] D. de Waard, M. van der Hulst, M. Hoedemaeker, and K. A. Brookhuis, "Driver behavior in an emergency situation in the automated highway system," *Transportation human factors*, vol. 1, no. 1, pp. 67–82, 1999.
- [20] J. Radlmayr, C. Gold, L. Lorenz, M. Farid, and K. Bengler, "How traffic situations and non-driving related tasks affect the take-over quality in highly automated driving," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 58, no. 1. Sage Publications, 2014, pp. 2063–2067.
- [21] M. Maltz and D. Shinar, "Imperfect in-vehicle collision avoidance warning systems can aid distracted drivers," *Transportation research part F: traffic psychology and behaviour*, vol. 10, pp. 345–357, 2007.
- [22] M. S. Young and N. A. Stanton, "Attention and automation: new perspectives on mental underload and performance," *Theoretical Issues in Ergonomics Science*, vol. 3, no. 2, pp. 178–194, 2002.
- [23] G. H. Walker, N. A. Stanton, and M. S. Young, "An on-road investigation of vehicle feedback and its role in driver cognition: Implications for cognitive ergonomics," *International Journal of Cognitive Ergonomics*, vol. 5, no. 4, pp. 421–444, 2001.
- [24] J. Yen, X. Fan, S. Sun, T. Hanratty, and J. Dumer, "Agents with shared mental models for enhancing team decision makings," *Decision Support Systems*, vol. 41, no. 3, pp. 634–653, 2006.
- [25] L. Takayama, D. Dooley, and W. Ju, "Expressing thought: improving robot readability with animation principles," in *Proceedings of the 6th international conference on Human-robot interaction*. ACM, 2011, pp. 69–76.
- [26] A. Steinfeld, "Interface lessons for fully and semi-autonomous mobile robots," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, vol. 3. IEEE, 2004, pp. 2752–2757.
- [27] K. Driggs-Campbell and R. Bajcsy, "Experimental design for human-in-the-loop driving simulations," Master's thesis, EECS Department, University of California, Berkeley, May 2015. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-59.html>
- [28] "PreScan Simulation Software," <https://www.tassinternational.com/prescan>, 2014, accessed: 2014-09-30.
- [29] "TkInter," <https://wiki.python.org/moin/TkInter>, accessed: 2016-05-01.
- [30] T. Kondoh, T. Yamamura, S. Kitazaki, N. Kuge, and E. R. Boer, "Identification of visual cues and quantification of drivers' perception of proximity risk to the lead vehicle in car-following situations," *Journal of Mechanical Systems for Transportation and Logistics*, vol. 1, no. 2, pp. 170–180, 2008.
- [31] D. Sadigh, K. Driggs Campbell, A. A. A. Puggelli, W. Li, V. Shia, R. Bajcsy, A. L. Sangiovanni-Vincentelli, S. S. Sastry, and S. A. Seshia, "Data-driven probabilistic modeling and verification of human driver behavior," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2013-197, Dec 2013. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-197.html>