TRANSFER FROM HIGHLY AUTOMATED TO MANUAL CONTROL: PERFORMANCE AND TRUST

Chris Schwarz
Timothy Brown
John Gaspar
National Advanced Driving Simulator, The University of Iowa, USA
Clara Keum
Department of Epidemiology, The University of Iowa, USA
Paper Number 17-0426

ABSTRACT

The development of automated vehicles is ongoing at a breakneck pace. The human factors challenges of designing safe automation systems are critical as the first several generations of automated vehicles are expected to be semi-autonomous, requiring frequent transfers of control between the driver and vehicle. A driving simulator study was performed with 20 participants to study transfers of control in highly automated vehicles. We observed driver performance and measured comfort as an indicator of the development of trust in the system. One study drive used a more capable automation system that was able to respond to most events by slowing or changing lanes on its own. The other study drive used a less capable automation system that issued takeover requests (TORs) in all cases. Thus there was a change in reliability over the course of the study drives; some participants experienced the more-capable system first followed by the less-capable system, and others had the opposite experience. We observed three types of comfort profiles over the course of the drives. Some drivers started out very comfortable, while others took a long time to become comfortable. Takeovers were split into physical takeover, visual attention, and vehicle stabilization. Response time and performance measures showed that there was a 15- to 25-second period between the physical takeover and a return to normal driving performance. This confirms some observations in previous studies on transfer of control.

INTRODUCTION

Automated vehicles are under active development by many auto manufacturers, tier 1 suppliers, and technology companies. The projected benefits of automated vehicles are many and varied, but so are the concerns over their technical limitations, legal barriers, and human factors challenges. The National Highway Traffic Safety Administration (NHTSA) started actively investigating automated vehicles in 2012 and has released their first policy document [1].

Transfer of control is a complex topic given the number of possible scenarios. An analysis of takeover types by Lu, et. al. [3] resulted in a unified framework that can be used to think about automation handoffs. A transfer of control (or takeover, transition, handoff) can result in the driver being in control (DC) or the automation being in control (AC). Moreover, they can be driver initiated (DI) or automation initiated (AI). This results in the four possible categories of transfers: DIDC, DIAC, AIDC, and AIAC. The underlying reason for the takeover can be classified as optional or mandatory. An optional transfer could be skipped with no adverse consequences, whereas missing a mandatory transfer would result in a safety critical event or crash.

This study was primarily concerned with automation level 3, termed *conditional automation* by SAE [2], in which the vehicle takes both longitudinal and lateral control. Whereas level 2 automation requires the operator to supervise the automation and scan the roadway for hazards, level 3 allows the operator to engage in other tasks, provided they can become available to take over again should the system request it. Both level 2 and level 3 raise concerns about how quickly the driver can take back control should they need to. Drivers can quickly become out of the loop and then have to regain situational awareness (SA) to effectively drive again. Because of this, our main interest was in studying mandatory AIDC transfers.

Bainbridge pointed out that humans are challenged when performing under time pressure and that when automation takes over the easy tasks from an operator, difficult tasks may become even more difficult [4]. Stanton and Marsden highlighted several potential problems that could plague automated vehicles, specifically when drivers must reclaim control from automation. These include over-reliance, misuse, confusion, reliability problems, skills maintenance, error-inducing designs, and shortfalls in expected benefits [5], [6]. The lack of situational awareness that occurs when a driver has dropped out of the control loop has been studied for some time in several different contexts [7]–[9].

More recently, it has been shown that drivers had significantly longer reaction times in responding to a critical event when they were in automation and required to intercede, compared to when they were driving manually [10]. More recent data suggest that drivers may take around 15 seconds to regain control from a high level of automation and up to 40 seconds to completely stabilize the vehicle control [11].

Takeover requests are issued by the automation to let the operator know that they should take back manual control of the dynamic driving task (DDT). The appropriate timing of such TORs has been a topic research recently. Takeover request timings of five and seven seconds ahead of encountering an obstacle in the road were tested in a driving simulator [12]. While it was possible for drivers to take over in only a couple of seconds in both conditions, there were more braking responses and less time to check their blind spots in the five-second timing condition. Some of the extra time in the seven-second condition was used for decision-making and was valuable for avoiding sudden braking responses.

A NHTSA-funded test track study used both imminent and staged TORs, where the imminent TOR was issued once with an external threat and once without [13]. The staged alert had four phases as follows: 1) a tone followed by an informational message, 2) a verbal alert with a cautionary message, 3) a repeated tone in addition to an orange visual alert, and 4) a repeated imminent tone with a red alert. The visual components were text messages with associated colors to indicate urgency. The four messages were the following:

1. Prepare for manual control
2. Please turn off autodrive
3. Turn off autodrive now (orange)
4. Turn off autodrive now (red)

The average response time to an imminent alert was 2.3 seconds without an external threat and 2.1 seconds with one. The average response time to the staged alert was 17 seconds, which may have been partly due to a countdown that accompanied the informational warning.

A driver's trust in automation greatly influences whether that automation is used appropriately, misused, or disused. Trust should be calibrated appropriately so that a driver does not over- or under-trust an automated system [14]. Lee and See proposed a closed-loop conceptual model of a dynamic process that governs trust, recognizing that trust might be considered as a function over time that can rise and fall.

Trust and comfort are correlated constructs that are both important for human-robot interaction [15], [16]. Indeed, it is hard to imagine the development of trust without some degree of comfort being present. Sanders et al. identified four factors of trust: performance, reliance, individual differences, and collaboration. Another breakdown of trust included the following factors: predictability, dependability, faith, and overall trust [17], [18].

A word on simulator fidelity is warranted. A series of driving simulator studies on adaptive cruise control done in the 1990s with and without motion showed similar results, and the authors concluded that motion may therefore not be necessary [19]. However, most recent driving simulation studies in vehicle automation have used higher-fidelity systems with motion bases. The 'feel' of the car from a simulator's motion cues is critical to a driver who may be completely visually disengaged from the driving task, as is the case in higher automation levels.

Objectives

This project was focused on transfers from conditional automation to manual control. The study events were mandatory takeovers that could be thought of as expected (approaching highway exit on route) and unexpected (approaching a slow-moving vehicle). The study was conducted using the NADS-1 high fidelity motion-base driving simulator, located at the University of Iowa.

The study was designed to address the following research questions:

1. To what degree do drivers trust the automation?
2. Does less-capable automation decrease trust, and how does reliability influence trust in automation?
3. When do drivers choose to begin an expected transfer of control, and how long does it take?
4. After manual takeovers, how long does it take for the driver to return control to the automation?
5. How long does an unexpected transfer of control take, including vehicle stabilization?
6. Does the act of taking manual control have any associated performance decrements?

It was expected that there would be decrements to the quality of the transfer due to the need to regain SA while at the same time assuming vehicle control. Moreover, it was also expected that automation failures, resulting in TORs, would damage the driver's trust in the system and that the effects of that reduced trust might be observed in subsequent driving and takeover choices.

We did not consider failures in the sense that the vehicle failed to issue a takeover request (TOR), which is a particularly concerning failure mode in its own right. Thus while the vehicle failed to navigate some study events, it always successfully issued TORs.

The term reliability was used in this research to indicate a change in the way the automation worked on similar events. In the more capable condition, the automation was able to navigate most study events by changing lanes. However, in the less capable condition it always issued a TOR. The automation capability condition was manipulated within-subjects across two drives, and the order of the drives was counterbalanced across drivers, resulting in a change in reliability for all drivers.

METHODOLOGY

A 2 (drive) x 2 (age) x 2 (gender) mixed design was used for this study. The within-subject independent variable was the automation reliability. The between-subject independent variables were gender (male, female) and age (18 - 25, 25 - 55). The age variable was blocked by using the minimization method to balance out the number of participants in each group. A total of 20 participants provided written informed consent and participated in the study.

Apparatus

The National Advanced Driving Simulator (NADS) is located at the University of Iowa. The NADS-1 simulator consists of a 24-foot dome in which an entire car cab is mounted. All participants drove the same vehicle—a 1996 Malibu sedan. The motion system, on which the dome sits, provides 400 square meters of horizontal and longitudinal travel and ±330 degrees of rotation. The driver feels acceleration, braking, and steering cues much as if he or she were actually driving a real vehicle. High-frequency road vibrations up to 40 Hz are reproduced from vibration actuators placed in each wheel well of the cab. A picture of the NADS-1 simulator and an image from the interior of the dome are shown in Figure 1.

The NADS-1 displays graphics by using 16 high-definition projectors that provide 360-degrees of horizontal, and 40-degrees of vertical, field of view. The NADS produces a thorough record of vehicle state (e.g., lane position) and driver inputs (e.g., steering wheel position), sampled at 240 Hz.

Figure 1. NADS-1 driving simulator (left) with a driving scene in the dome (right).

The cab is equipped with a Face Lab™ 5.0 eye-tracking system that is mounted on the dash in front of the driver's seat above the steering wheel. In the best-case scenario, where the head is motionless and both eyes are visible, a fixated gaze may be measured with an error of about 2º. With the worst-case head pose, accuracy is estimated to be about 5º. The eye tracker samples at a rate of 60 Hz.

Driving Scenarios

Participants completed a seven-minute practice drive followed by two thirty-minute study drives containing the same set of events (see Table 1). The study drives involved typical vehicle control in a variety of situations. Once the driver achieved highway speed, he or she was instructed to engage the automation by pressing a button on the steering wheel.

Table 1. Scenario events in the more and less capable drives (A and B) with varying takeover request (TOR) timing.

Event	More Capable (A)	Less Capable (B)
#1 Work zone	No TOR	10 sec. TOR
#2 Missing lane lines	No TOR	10 sec. TOR
#3 Sharp curve	No TOR	10 sec. TOR
#4 Slow lead vehicle	10 sec. TOR	5 sec. TOR
#5 Exit highway	30 sec. TOR	30 sec. TOR

The practice drive scenario served to adapt participants to driving in the simulator, as well as expose them to automation control transfers and TORs. All five events existed in both study drives, but in different orders and with different automation capabilities. Moreover, the locations of the events as well as the starting and ending locations of the drives were also varied to minimize predictability. Towards the end of each drive, an expected takeover request took place before a scheduled exit off the highway. The five main events are summarized in Table 1.

Occasionally, a lead vehicle would slow from the speed limit to 55 mph for a short time, forcing the automation to slow the participant's vehicle as well. Then the lead vehicle sped back up to the speed limit. These brief disturbances drew the operator's attention and provided experiences in which the automation behaved as desired with no loss of capability. It was expected that these instances would help to build trust in the system.

Driver Vehicle Interface

Automated driving was indicated by a visual icon on a high heads-up display. Takeover requests were composed of both visual and audio cues. Visual cues appeared on the same display. When the driver needed to transfer control, a chime sound played with the appearance of a visual sign saying to either turn on or off the automation. Depending on each event and scenario, a TOR took place either 5, 10, or 30 seconds prior to the event. If the driver did not transfer control from automated to manual in some set interval after the TOR fired, the automation system slowed the vehicle down and pulled over to the side of the highway. This fallback strategy is characteristic of SAE Level 4 automation, though participants were not trained on it ahead of time, and it was never encountered in the study. All four possible display icons are shown in Figure 2.

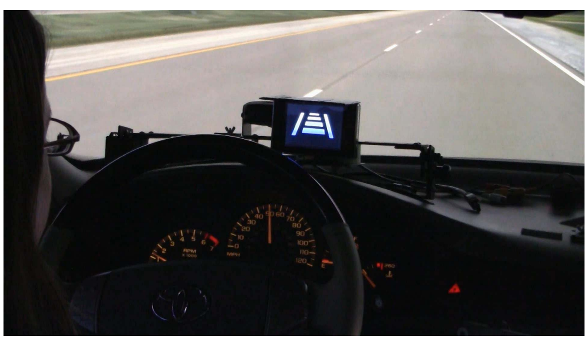

(a)

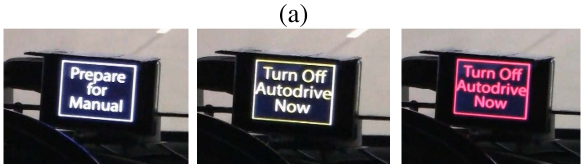

(b) (c) (d)

Figure 2. Automation interface in high heads-up display location: (a) automated-mode icon in blue, (b) informational warning in white, (c) cautionary alert in yellow, (d) imminent alert in red.

Non-Driving Task

Participants were asked to work on trivia questions from the website Trivia Plaza (www.triviaplaza.com) as an alternative primary task while the vehicle was under automated control during both drives. Trivia Plaza is a website that offers numerous sets of questions in nine major categories (see Figure 3). Within each category, there are many subcategories (e.g., subcategories of "Movie" include various time periods, genres, production companies, etc.). The intent was to provide a task that all participants could be equally engaged with, by finding topics of greatest interest to them.

An iPad was given to each participant for the duration of the drives to allow access to the website. In order to encourage participants to be actively involved in trivia, they were told to pick any topic(s) that they were interested in and that any participant who reached a cumulative score of 100 or higher would receive a bonus compensation of $15. Participants could play multiple times to reach the given score. In reality, all subjects received the $15 bonus.

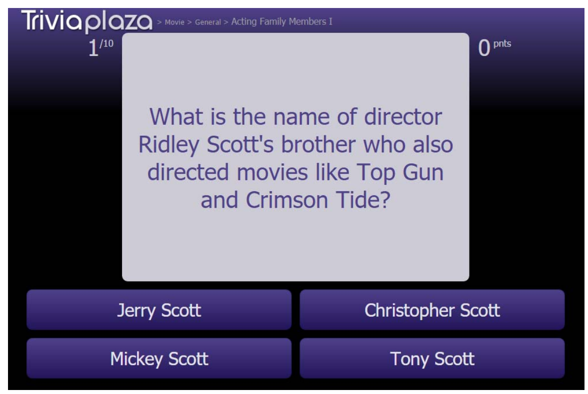

Figure 3. Example screen from Trivia Plaza (www.triviaplaza.com).

Driver Comfort

The amount of comfort an operator had in the automation during their drives was probed at semi-regular intervals using an online survey that appeared on a display located in front of the cab's center console. The single question asked the operator to rate his or her level of comfort at that moment on a scale of 1 (Very Comfortable) to 7 (Very Uncomfortable). The wording of comfort was selected as an overall approximation of the more complex concept of trust and was thought to estimate the participants' nascent level of trust in a system that was new to them.

Two such comfort probe surveys were administered in the practice drive. There were eight additional surveys in each main drive, for a total of 18 comfort measurements. They were spaced in between events, and nothing related to any event was happening at the time the surveys were administered. Sometimes the survey occurred after one of the four main events, but sometimes it occurred after the 'filler' event during which a lead vehicle slowed momentarily.

Dependent Measures

Data was collected from three main sources. Simulator data files contained many variables, including driver inputs and vehicle signals. Eye tracker data was recorded to log files from the FaceLab system. Lastly, post-drive surveys were administered to collect additional data on comfort and attitudes towards automated vehicles, and a comfort probe survey was given at semi-regular intervals in the cab during the study drives. The simulator and eye tracker data were processed using a data reduction script in Matlab to obtain several dependent measures used in the analysis.

Two types of measures were calculated. The first set was calculated once per event and is listed in Table 2. These measures included response times, eye gaze, and information about the use of automation (see Table 2). The percent road center (PRC) gaze [20] measured the percentage of time that the driver's gaze was directed at the front scene, computed in a running 17-second window [21].

Table 2. Dependent measures, calculated once per event.

Measure	Description
PctAuto	Percentage of event time spent in automated mode
TakeOverRT	Response time to take over from automation after warning
GiveBackRT	Response time to give back control to automation after cue
MeanPrc17Auto	Average PRC gaze while in automated mode
MedPrc17Auto	Median PRC gaze while in automated mode
MeanPrc17Manual	Average PRC gaze while in manual mode
MedPrc17Manual	Median PRC gaze while in manual mode
DurationManual	The time that was spent in manual mode
Manual	Did the driver take back control from the automation?

A second type of dependent measure was recorded at regular intervals either after the beginning of manual driving mode, or after the end of manual driving mode in the event. A fixed interval spacing of five seconds was used, and up to 12 segments, or one minute, were computed. These measures created a type of longitudinal, or time sequence, data that could be analyzed for trends. The approach was adapted from the methodology used by Merat el al. [11]. The longitudinal dependent measures are summarized in Table 3.

Table 3. Longitudinal dependent measures, calculated in five second segments.

Measure	Description
MinSpeed	The minimum speed in each manual segment (mph)
MeanSpeed	The average speed in each manual segment (mph)
SR	The steering reversal rate in each manual segment, calculated in a 15-second running window (rev/sec)
SDLP	Average value of standard deviation of lane position in each manual segment, calculated in a 15-second running window (ft)
HFSteer	High-frequency steering content in each manual segment
PRC	Percent road center gaze in each manual segment, calculated in a 17-second running window (%)
PRCpost	Percent road center gaze in each segment after return to automated mode, calculated in a 17-second running window (%)

The steering reversals and high-frequency steering (HFSteer) measures were also adapted from the methodology in [11]. Steering reversals count the number of one-degree reversals in a time period. The steering reversal rate per second was then calculated by dividing by the number of seconds in the segment. The HFSteer measure is based on a high-frequency control of steering computation that is defined as the ratio between the power of a high-frequency band of steering activity to the power of a lower-frequency band [22], [23].

RESULTS

Results on Operator Trust

<u>How much did operators trust the automation?</u>

The R statistical software language [24] was used to analyze the simulator and eye tracker measures. Box Cox transformations were applied to the dependent measure, where appropriate, to optimize the normality of the residual error. Normality was tested by observing the Q-Q plot of the residuals as well as by running a Shapiro-Wilk test to see if the null hypothesis of normality should be rejected. Additionally, a cluster analysis was used to identify three distinct profiles of longitudinal comfort that were observed among the participants.

The log of the in-cab comfort score was used as the main trust measure. All 18 measurements in a drive constituted a longitudinal comfort profile that evolved in ways unique to each individual. Each participant's longitudinal comfort profile is plotted individually in Figure 4. The scenario is coded both by color and by marker shape. No significant effect of age, gender, or order of the drives was found on the development of comfort, tested using growth curve models with linear piecewise time segments.

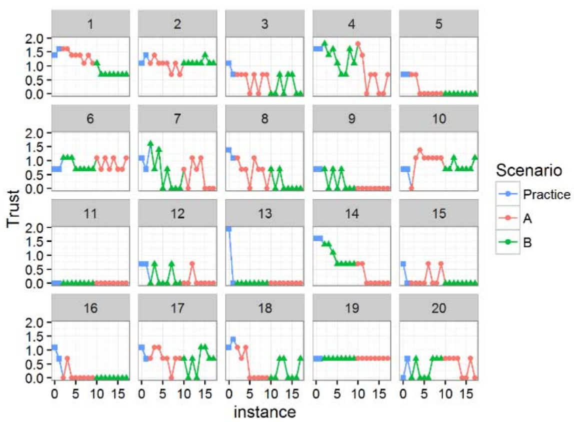

Figure 4. Longitudinal comfort (log of 18 comfort responses) for all participants across three study drives. Drive A used more capable automation, while drive B used less capable.

A hierarchical clustering analysis was conducted using the random intercept and two random slopes from the growth curve model. Three clusters were selected from the analysis and participants were assigned to one of the three. Figure 5 shows the longitudinal comfort profiles once again, this time with 95% confidence intervals from the random effects overlaid on each plot. Additionally, the cluster for each participant is color-coded in the figure.

The three clusters may be easily described on inspection of Figure 5. The participants in cluster one gradually increased in comfort (the log of the response is inversely proportional to comfort) over the course of the practice drive and two main drives. Participants in cluster two started with about the level of comfort that they maintained throughout their three drives. Finally, participants in cluster three started with less comfort, but became more comfortable over a fixed amount of time and then leveled off for the remainder of the drives. Participant 13 may be an outlier if the first large comfort response was an aberration. Participant 4 was unusual in that the responses indicated a loss of comfort near the end of the first drive (identified as Drive B, or the less-capable automation system, from Figure 4).

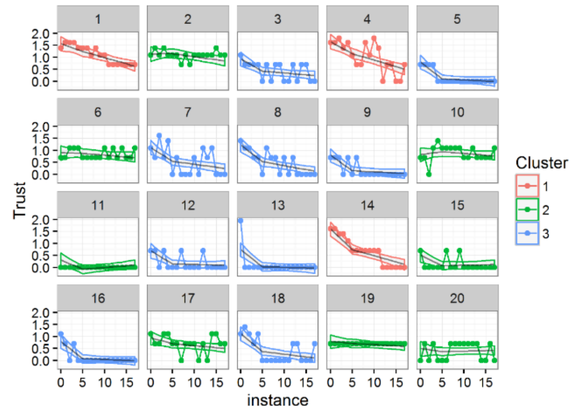

Figure 5. Three comfort profile clusters. Ribbon overlays show 95% confidence interval of the random effects model fit.

<u>How did participants rate their trust retrospectively?</u> For the retrospective trust survey data, the restricted range and ordinal scale of the data associated with Likert-type survey responses required that care be taken in that analysis. Although there is significant debate over the acceptability of various analysis approaches and whether these data can be considered as interval scale and analyzed with ANOVA, Sullivan and Artino [25] present an argument that ANOVA is an appropriate technique. Accordingly, the SAS general linear model (GLM) procedure was used to conduct an ANOVA on the post-drive survey data. Scenario (more or less capable (A or B)) and order (first or second drive (1 or 2)) were treated as within-subjects factors for each of four questions where participants provided comfort responses.

The first question asked participants to indicate how comfortable they felt when transferring into automated mode. Overall, participants felt quite comfortable, and there were no significant effects or interactions involving either scenario or order ($p > 0.05$). The second question asked participants how comfortable they felt when resuming manual control back from the automation. The main effect of order was marginally significant ($p = 0.09$), suggesting that drivers tended to be less comfortable in their first drive relative to their second drive. This is to be expected as drivers grew more familiar with the automation and transferring control.

The third question asked drivers how comfortable they felt when the automation failed and they had to regain control. Neither the main effect of order nor scenario reached significance, nor did the order by scenario interaction ($p > 0.05$). The final Likert scale question asked participants how comfortable they felt when driving in automated mode. Again, the main effects of order and scenario and the order by scenario interaction did not reach significance ($p > 0.05$).

These results generally suggest that the capability of the automation (scenario) and the order in which drivers experienced the different conditions had a limited effect on drivers' retrospective perceptions of comfort in interacting with the automation.

Results on Simulator Measures

How long to transfers of control take? Transfers of control from automated to manual operation have several phases that should be considered individually, though some are more difficult to study than others. Situational awareness, for example, is a difficult concept to define, much less measure, and we do not attempt it here, though visual attention is likely a good minimum bound on the time required to regain it. Four phases of takeover from automation are presented in Table 4. Note that order is not implied in the table, as SA could be fully regained before the physical takeover is initiated.

Table 4. Phases of takeover from automated to manual mode

Takeover Phase	Dependent Measure
Physically taking control by pressing the transfer button or the brake pedal	Takeover response time from cautionary TOR
Physically stabilizing control of the vehicle after taking control	Longitudinal dependent measures for steering and lane keeping
Visually attending to the dynamic driving task	Longitudinal dependent measure for PRC gaze during manual mode
Regaining full situational awareness	None

The physical takeover phase may be characterized by the drivers' response times in returning to manual mode after being given a TOR. Events 1 through 4 used cautionary TORs. The average response time was 4.13 seconds with a standard deviation of 1.04 seconds (see Figure 6a). The exit event, event 5, first issued an information TOR, followed by a cautionary TOR and an imminent TOR, each lasting for 10 seconds.

Observe in Figure 6b that the distribution of response times for event 5 is tri-modal. Some people responded after the first stage TOR and some after the third one. One person responded after 30 seconds. The first group had a mean time of 7.60 seconds with standard deviation of 1.28 seconds. The middle, largest, group had a mean response time of 22.37 seconds with standard deviation of 0.85 seconds. The participant in the third group responded at 31.57 seconds. Three-way ANOVAs were run on takeover response time for each event using order, gender, and age. No significant effects of these conditions were found.

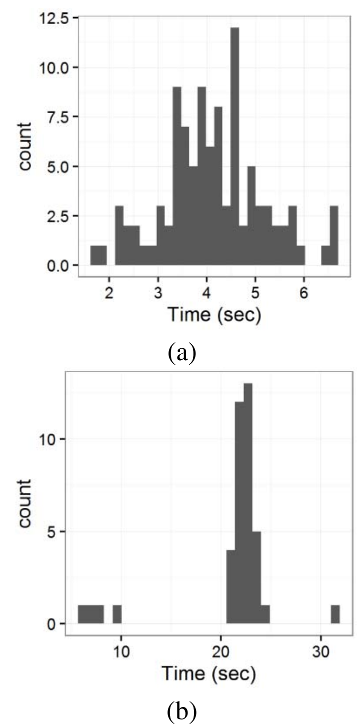

Figure 6. Distribution of response time to take back manual control after a TOR for (a) events 1 through 4, and (b) event 5.

The third phase of manual takeovers considers the time required for the driver to become fully visually engaged in the dynamic driving task. We used the percent road center (PRC) gaze measure recorded using the eye tracker to indicate visual attention. Percent road center has been used not only as a measure of visual distraction, but also to detect cognitive distraction. Simply put, PRC has a normal range, and values that are too low or too high indicate a lack of proper attention.

After manual takeovers, PRC gaze increased as drivers returned their gaze to the road until achieving normal gaze patterns once more. The PRC gaze was calculated on a 17-second running window, which has been used for the detection of distraction [21]. The increasing piece of the PRC gaze trend, up until it peaked, was fit to a linear model, and linear interpolation (or extrapolation, as appropriate) was used to estimate the time at which the PRC would reach 0.7. The distribution of these times is shown in Figure 7. In actuality, the PRC never reached 0.7 in some events for some participants. Such cases caused the increasing trend to have a very shallow slope, resulting in very large estimates for the 0.7 intercept time. Nevertheless, the estimate is useful as a way to compare events and participants against one another.

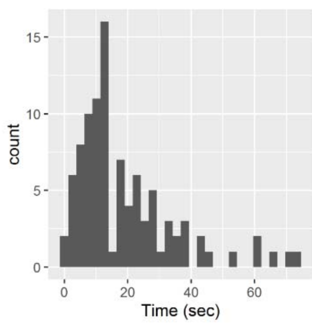

Figure 7. Distribution of times projected for PRC to reach 0.7 after transfer to manual mode

Transfers of control from manual to automated mode are simpler in that stabilization and situational awareness are not factors after the transfer. Rather, analyzing transfers to automated mode may tell us about the degree of trust the operator has in the automation. After each event, an audio/visual cue was given to the driver that they could once again transfer control to the automation. The response time was measured from the time this cue was issued. The distribution of response times for the driver to hand back control to the automation is shown in Figure 8. After removing the times larger than 20 seconds as outliers, the mean response time was calculated to be 5.31 seconds with standard deviation of 3.15 seconds.

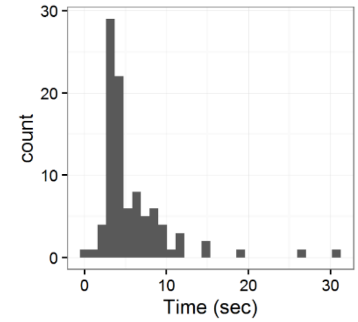

Figure 8. Distribution of response time to give back control to the automation after a reminder cue in events 1 through 4.

After control was returned to the automation, the PRC gaze dropped until the driver engaged once more with the trivia task. The PRC gaze trend was fit to a linear model, and the time was estimated at which the PRC would reach 0.1. A distribution of these times is shown in Figure 9.

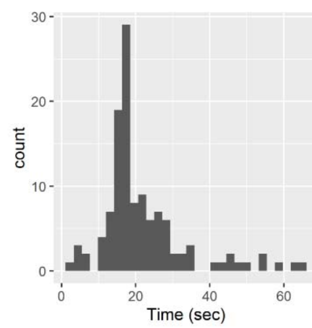

Figure 9. Distribution of times for PRC to reach 0.1 after transfer to automated model.

Were there performance decrements after manual takeovers? The second phase of manual takeover includes the time required to stabilize physical control of the vehicle. The high-frequency control of steering, captured in the HFSteer measure, is thought to be sensitive to distraction. A larger amount of variance was observed in the HFSteer measure in the first six time segments, while less variance was observed in the last six time segments. This is shown in Figure 10.

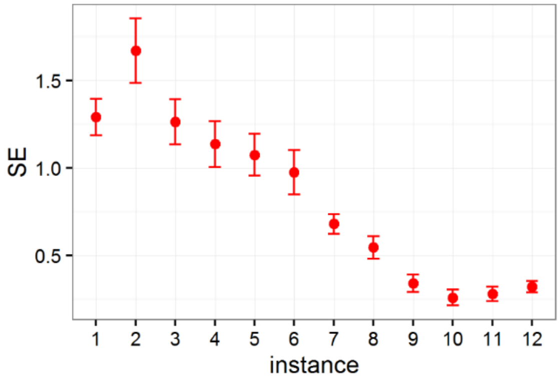

Figure 10. Standard error of HFSteer measure across all participants and all events for each time segment after a manual takeover

CONCLUSIONS

Twenty participants took part in an automated driving study using the NADS-1 motion base driving simulator. The automation was described generally as SAE Level 3 (conditional automation), however it was implemented as SAE Level 4 with a fallback mode to mitigate the risk that an automated vehicle would actually collide with a lead vehicle or drive through a work zone. Those negative outcomes did not happen, and the fallback mode was not needed in any of the events.

Comfort was measured using an online probe survey that was administered twice during the practice drive and eight times during each main drive. Also, a post-drive survey was administered after each main drive; it asked the participants to retrospectively consider their comfort with the automation. We surmised that asking about comfort would be an effective way to capture the nascent trust of an operator just becoming familiar with an automation system. Future work could delve deeper into multiple facets of trust, including predictability/performance, dependability/reliance, faith, and collaboration.

A cluster analysis revealed three distinct longitudinal comfort profiles from the probe surveys. One cluster started with a high level of comfort and stayed that way. Another started with a lower level of comfort, but it gradually increased after a few surveys and then stayed level. A third cluster started with low comfort and gradually increased over the course of the practice and two main drives. Apart from single instances of reduced comfort, only participant 4 showed a temporary trend of decreasing comfort. We could not associate the clustering with age, gender, or order. It may be that it is associated with some latent variable such as sensation seeking or a personality trait. The longitudinal comfort profiles support the notion that trust can be modeled as a function of time, especially in the sense that instantaneous levels of trust depend on their previously measured levels [31].

The physical response times to TORs and automation reminders were both under 10 seconds (4.13 sec +- 1.04 sec and 5.31 sec +- 3.15 sec, respectively). Visual attention to the driving task was measured using the percent road center gaze, calculated over a 17-second running window. There were many instances in which it took a driver more than 20 seconds to return to normal forward gaze after a transfer.

Consideration of the response times for physical takeovers, stabilization, and visual attention leads to concern for the driver's safety after taking control. Drivers are capable of physically taking over control in less than five seconds. However, PRC gaze showed that it could take 20 seconds or more to return their full attention to the roadway. Additionally, the variation in high-frequency steering offers evidence that drivers do not return to their normal driving control for up to 30 seconds. These results imply there could be a 15- to 25-second gap during which the driver may be vulnerable to missing a response to a safety-critical event at an inopportune moment.

The main limitations of this study were that it used a fairly small sample size (20 participants), and that it was not able to fully explore the different dimensions of trust. Future research should address both of those limitations. Additionally, the inclusion of safety-critical events and latent hazards, would allow a better judgement of whether the driver has regained SA and whether the takeover times observed have an adverse effect on safety. We modeled our driver-vehicle interface (DVI) largely on previous research. However, there is still much that could be done to test different modalities and timing for DVI design. Finally, we conjecture that the best DVI would be one that is capable of monitoring the driver and adapting elements of the DVI, transfers of control, and other aspects of the automation to the perceived state of the operator.

REFERENCES

[1] NHTSA, "Federal Automated Vehicles Policy: Accelerating the Next Revolution in Roadway Safety." NHTSA, Sep-2016.

[2] SAE, "Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated

Driving Systems," SAE, Standard J3016, Jan. 2014.
[3] Z. Lu, R. Happee, C. D. D. Cabrall, M. Kyriakidis, and J. C. F. de Winter, "Human factors of transitions in automated driving: A general framework and literature survey," *Transp. Res. Part F Traffic Psychol. Behav.*, vol. 43, pp. 183–198, Nov. 2016.
[4] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, Nov. 1983.
[5] N. A. Stanton and P. Marsden, "From fly-by-wire to drive-by-wire: Safety implications of automation in vehicles," *Saf. Sci.*, vol. 24, no. 1, pp. 35–49, Oct. 1996.
[6] N. A. Stanton, M. Young, and B. McCaulder, "Drive-by-wire: The case of driver workload and reclaiming control with adaptive cruise control," *Saf. Sci.*, vol. 27, no. 2–3, pp. 149–159, Nov. 1997.
[7] M. R. Endsley, "Automation and Situation Awareness," in *Automation and human performance: Theory and applications*, vol. xx, Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc, 1996, pp. 163–181.
[8] M. R. Endsley and E. O. Kiris, "The Out-of-the-Loop Performance Problem and Level of Control in Automation," *Hum. Factors J. Hum. Factors Ergon. Soc.*, vol. 37, no. 2, pp. 381–394, Jun. 1995.
[9] D. B. Kaber and M. R. Endsley, "Out-of-the-loop performance problems and the use of intermediate levels of automation for improved control system functioning and safety," *Process Saf. Prog.*, vol. 16, no. 3, pp. 126–131, 1997.
[10] N. Merat and A. H. Jamson, "How Do Drivers Behave in a Highly Automated Car?," in *Proceedings of the 5th International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design*, 2009.
[11] N. Merat, A. H. Jamson, F. C. H. Lai, M. Daly, and O. M. J. Carsten, "Transition to manual: Driver behaviour when resuming control from a highly automated vehicle," *Transp. Res. Part F Traffic Psychol. Behav.*, vol. 27, Part B, pp. 274–282, Nov. 2014.
[12] C. Gold, D. Damböck, L. Lorenz, and K. Bengler, "'Take over!' How long does it take to get the driver back into the loop?," *Proc. Hum. Factors Ergon. Soc. Annu. Meet.*, vol. 57, no. 1, pp. 1938–1942, Sep. 2013.
[13] M. Blanco *et al.*, "Human Factors Evaluation of Level 2 and Level 3 Automated Driving Concepts," NHTSA, Washington, D.C., Final Report DOT HS 812 182, Aug. 2015.
[14] J. D. Lee and K. A. See, "Trust in Automation: Designing for Appropriate Reliance," *Hum. Factors J. Hum. Factors Ergon. Soc.*, vol. 46, no. 1, pp. 50–80, Mar. 2004.
[15] V. Evers, H. Maldonado, T. Brodecki, and P. Hinds, "Relational vs. group self-construal: Untangling the role of national culture in HRI," in *2008 3rd ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2008, pp. 255–262.
[16] T. Sanders, K. E. Oleson, D. R. Billings, J. Y. C. Chen, and P. A. Hancock, "A Model of Human-Robot Trust Theoretical Model Development," *Proc. Hum. Factors Ergon. Soc. Annu. Meet.*, vol. 55, no. 1, pp. 1432–1436, Sep. 2011.
[17] M. A. Price, V. Venkatraman, M. Gibson, J. D. Lee, and B. Mutlu, "Psychophysics of Trust in Vehicle Control Algorithms," in *Proceedings of the SAE World Congress*, Detroit, MI, 2016.
[18] B. M. Muir and N. Moray, "Trust in automation. Part II. Experimental studies of trust and human intervention in a process control simulation," *Ergonomics*, vol. 39, no. 3, pp. 429–460, Mar. 1996.
[19] N. A. Stanton and M. S. Young, "Vehicle automation and driving performance," *Ergonomics*, vol. 41, no. 7, pp. 1014–1028, 1998.
[20] T. W. Victor, J. L. Harbluk, and J. A. Engström, "Sensitivity of eye-movement measures to in-vehicle task difficulty," *Transp. Res. Part F Traffic Psychol. Behav.*, vol. 8, no. 2, pp. 167–190, Mar. 2005.
[21] J. D. Lee *et al.*, "Distraction Detection and Mitigation Through Driver Feedback: Appendices," Final Report DOT HS 811 547B, May 2013.
[22] J. R. Mclean and E. R. Hoffmann, "Analysis of Drivers' Control Movements," *Hum. Factors J. Hum. Factors Ergon. Soc.*, vol. 13, no. 5, pp. 407–418, 1971.
[23] J. R. McLean and E. R. Hoffmann, "The Effects of Restricted Preview on Driver Steering Control and Performance," *Hum. Factors J. Hum. Factors Ergon. Soc.*, vol. 15, no. 4, pp. 421–430, Aug. 1973.
[24] R Development Core Team, *{R: A language and environment for statistical computing}*. Vienna, Austria: R Foundation for Statistical Computing, 2009.
[25] G. M. Sullivan and A. R. Artino, "Analyzing and Interpreting Data From Likert-Type

Scales," *J. Grad. Med. Educ.*, vol. 5, no. 4, pp. 541–542, Dec. 2013.
[26] "Fitting Linear Mixed-Effects Models Using lme4 | Bates | Journal of Statistical Software." [Online]. Available: https://www.jstatsoft.org/article/view/v067i01. [Accessed: 28-Jul-2016].
[27] D. Mirman, J. A. Dixon, and J. S. Magnuson, "Statistical and computational models of the visual world paradigm: Growth curves and individual differences," *J. Mem. Lang.*, vol. 59, no. 4, pp. 475–494, Nov. 2008.
[28] K. P. Burnham and D. R. Anderson, "Multimodel Inference Understanding AIC and BIC in Model Selection," *Sociol. Methods Res.*, vol. 33, no. 2, pp. 261–304, Nov. 2004.
[29] S. Nakagawa and H. Schielzeth, "A general and simple method for obtaining R2 from generalized linear mixed-effects models," *Methods Ecol. Evol.*, vol. 4, no. 2, pp. 133–142, Feb. 2013.
[30] P. C. D. Johnson, "Extension of Nakagawa & Schielzeth's R2GLMM to random slopes models," *Methods Ecol. Evol.*, vol. 5, no. 9, pp. 944–946, Sep. 2014.
[31] J. D. Lee and N. Moray, "Trust, control strategies and allocation of function in human-machine systems," *Ergonomics*, vol. 35, no. 10, pp. 1243–1270, 1992.