Adapting Automated Vehicle Behavior to User Trust: a Driving Simulator Study

Julius Hörsting
University of Twente

Master Thesis: Faculty of Cognitive Psychology and Ergonomics (CPE)

First supervisor: Prof. Dr. Ing. Willem B. Verwey

Second supervisor: Dr. Simone Borsci

External supervisor: Dr. Francesco Walker

July 4, 2022

Table of Contents

Abstract

Automated vehicles are claimed to be the solution for a lot of issues that come with manual driving. They are promised to make driving safer, allow for engaging in secondary tasks, and be better for the environment. However, for automated vehicles to be adopted by the public, people need to have an adequate level of trust in them, meaning it should neither be too high nor too low. This study set out to investigate whether adapting a vehicle's speed to the user's trust leads to higher trust levels in individuals that tend to have low trust in automated vehicles. For that, in a between-subjects design, 45 participants were divided into 3 groups (one experimental and two control groups). Participants were asked to take a 15-minute ride in an automated vehicle in a simulator while continuously indicating their trust in the automated vehicle on a slider. Additionally, their electrodermal activity during the drive and their trust level before and after the drive were recorded. The results were unable to confirm that adapting the vehicle speed to the user's trust leads to higher trust levels compared to the two control groups. Further follow-up studies, with their design based on the findings as well as recommendations in this study, should be conducted to investigate adaptive speed changes further.

Keywords: automated vehicles, adaptive speed, trust in automation, trust calibration

Introduction

With about 1.3 million traffic-related deaths each year, about 20 to 50 million non-fatal accidents, and with traffic accidents being the leading cause of death among 5-29-year-olds, there is a strong need to make driving safer (World Health Organization, 2021). Therefore, the United Nations General Assembly's goal is to reduce the number of deaths and injuries on the road by 50% by 2030 (United Nations General Assembly, 2020). Automated vehicles (AVs) promise to make driving safer as well as improve overall road safety. AVs are claimed to provide the ability to support drivers when they experience risky or demanding driving scenarios (Sun et al., 2020). Those who no longer feel comfortable with driving manually or simply are unable to drive, could benefit from automated vehicles as they would (re-)gain the ability to move around on their own with a car (Sun et al., 2020; Walker et al., 2019).

However, while automated vehicles provide many benefits, there are quite a few challenges, with their adoption being one of the biggest (Ekman et al., 2018; Kaur & Rampersad, 2018; Zhang et al., 2019). A factor that plays a crucial role in the adoption of automated vehicles is trust (Ekman et al., 2018; Ekman et al., 2019; Lee & See, 2004; Sun et al., 2020). Trust leads to acceptance, which in turn is needed for the adoption and usage of AVs (Ekman et al., 2018). Acceptance is defined as "the attitude towards, or the willingness for use (or non-use), that an individual has of an advanced system" (Kaye et al., 2021, p. 2). Since both the driver and the AV rely on each other and using an AV always comes with some unpredictability, trust in AVs is vital (Ekman et al., 2018). It is crucial that drivers appropriately trust the AV in a way that the trust of the driver aligns with their expectations of the vehicle's capabilities (Sun et al., 2020). This is important as only with an adequately calibrated amount of trust will the benefits that an AV promises come into effect.

Transitioning from manual to automated driving requires allowing the vehicle to take control and trusting the AV to execute all driving-related tasks reliably and harmlessly (Sun et al., 2020). This transition to a more supervisory role when riding in an AV might result in decreased trust in the vehicle and its capabilities (Pettersson & Karlsson, 2015). Personalizing certain aspects of an AV's driving style (such as vehicle speed) might be one way of increasing user trust in AVs (Sun et al., 2020). While there are some studies that investigated different driving styles (e.g. Ekman et al., 2019), there is little, if any, research on the effect of personalized adaptive driving speeds on trust. Hence, the research reported in this paper aims to test whether adapting an AV's speed to user trust increases the trust in the driving capabilities of the AV.

Background

5 levels of automation

AVs can have a wide array of capabilities and features, with different levels of automation ranging from no automation (in which the driver is fully responsible for driving) to full automation (where automated systems, instead of the driver, are fully responsible for driving). A widely adopted classification is the one developed by the Society of Automotive Engineers (SAE) which has classified 5 levels of automation (see Figure 1). As this paper focused on Level 4 automated vehicles, explaining all levels of automation in detail is beyond the scope of this manuscript. For a full review of the 5 levels of automation, see SAE (2021a).

Figure 1

The 5 levels of automation as classified by the SAE (SAE, 2021b)

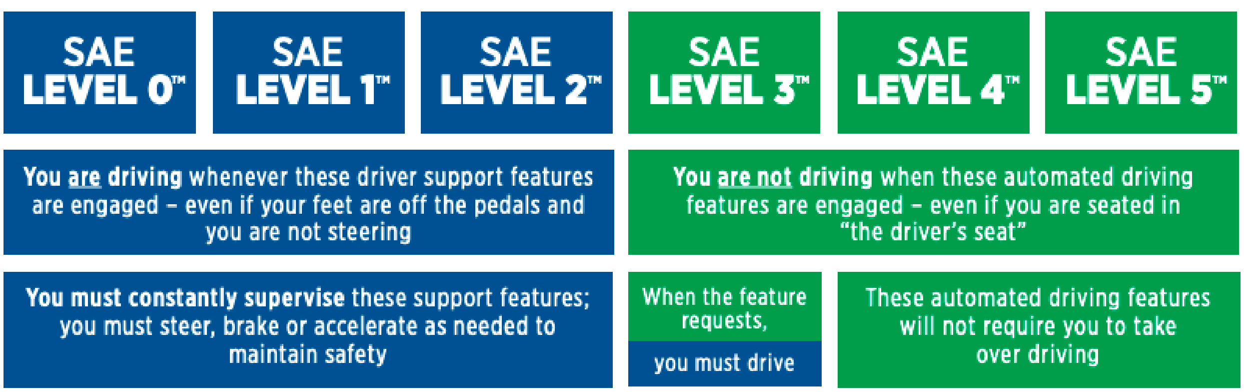

In brief, Levels 0 (no automation) to 2 (partial automation) require the driver to drive and monitor the vehicle (SAE, 2021a). Vehicles with these levels of automation include features and capabilities such as automatic braking, adaptive cruise control, and lane centering (SAE, 2021b). Some current commercial vehicles already use Level 2 automation systems (e.g. Tesla Autopilot, Volvo Pilot Assist, and Mercedes-Benz Drive Pilot) (Walker, 2021).

Level 3 (conditional automation) automated features only require the driver to take control of the car when the automated system is confronted with situations or circumstances, which are beyond the system's capabilities (SAE, 2021a). With Level 4 (high automation) automated features, the autonomous vehicle is fully automated, but only when in its operational design domain. The operational design domain of AVs are specific areas and/or circumstances in which the autonomous vehicle is able to operate fully automated. When leaving their operational design domain, the vehicle will ask the driver to take over the driving and the vehicle can be driven manually. Should the driver not comply with the takeover request, the vehicle will switch into a minimal risk condition (the vehicle will stop with the lowest risk possible). An example of a Level 4 automated feature would be a local driverless taxi that can operate fully driverless in that local area only (SAE, 2021b). For Level 5 (full automation) automated features, the driver does not need to take over the driving task at all and the automated vehicle will be able to drive fully automatically in all situations and circumstances. While there is no Level 3 or above automated vehicle available on the public market yet, Mercedes-Benz is expected to release a Level 3 system for their S-Class in 2022 (Reyes, 2021).

Trust in automated driving

Lee and See (2004) define trust as "the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability" (p.51). Trust influences to what degree people will accept and use an automated system as well as the way in which it will be used (Parasuraman & Riley, 1997; Walker et al., 2019). It plays a crucial part in overcoming situations in which users are confronted with uncertainty and cognitive complexity that they were, for example, not used to in manual driving (Lee & See, 2004).

Promoting an adequate level of trust in AVs is pivotal to making the best and safest use of an AV's capabilities (Hoff & Bashir, 2014; Lee & See, 2004). Users may sometimes trust the automation of the vehicle too much, whereas at other times they may trust the automation too little. This has been termed by Parasuraman and Riley (1997) as misuse and disuse, respectively. Drivers' misuse or disuse as well as the degree to which people accept AVs is based on their initial trust (Lee & See, 2004; Walker et al., 2019). Additionally, it depends on the driver's previous experience and how reliable they perceive the AV to be.

Previous studies have shown that drivers' trust in an AV is influenced by the driving style of the vehicle (Ekman et al., 2019). However, the standard driving style settings of an AV (as set by engineers that programmed the vehicle) might not be the most suitable or most comfortable for a lot of people (Trende et al., 2019). This was also found in a recent study by Walker et al. (2020), where they found that the perception of engineers and users concerning an AV's reliability often does not match. Mismatched and uncomfortable driving experiences can prevent the adoption of AVs due to experiencing increased amounts of stress and workload (Trende et al., 2019). It is therefore important that the capabilities of the car match the expectations, and thereby the trust, of the user.

Electrodermal activity

Nowadays, trust is often measured by questionnaires, such as the Jian, Bisantz, and Drury (2000) trust scale (Walker et al., 2019). However, these questionnaires do not offer insights into any real-time fluctuation in trust, such as when users encounter specific situations or conditions (Azevedo-Sa et al., 2020; Walker et al., 2019). A way of collecting real-time data of trust is to register changes in skin conductance. Depending on the difficulty of an encountered situation changes in skin conductance occur, with low skin conductance indicating a less challenging situation and high skin conductance reflecting a more challenging situation (Erath et al., 2009, as cited in Song et al., 2019). These changes in skin conductivity are called electrodermal activity (EDA) and happen outside of the user's conscious awareness; the user cannot actively control these changes (Walker et al., 2019). Several studies have examined the correlation between EDA and trust and have found that EDA can indeed give a good indication of a user's trust (Akash et al., 2018; Khawaji et al., 2015). While the aforementioned studies have investigated EDA outside of the domain of AVs, trust also seems to be influencing the physiological state of a user while driving an automated vehicle (Morris et al., 2017; Walker et al., 2019). Walker et al. (2019) found that EDA can give a real-time indication of the driver's trust while driving. A high electrodermal activity, as compared to a baseline, has been found to indicate a low level of trust while being driven in an AV (Morris et al., 2017; Walker et al., 2019). Even though EDA can provide real-time indications of user trust, it seems hardly useful for AVs in practice (Azevedo-Sa et al., 2020). Nevertheless, using EDA in a research setting to measure driver trust in AVs might provide crucial insights needed for the development of future automated systems.

Adapting the speed of automated vehicles to user trust

To this day, there is little to no insight into the impact that personalization of the AV's driving style might have on trust in AVs and at the moment AVs seem to be developed to only fit the standard driver (Sun et al., 2020). Currently, manufacturers fail to account for the differences in preferred driving styles between people. Not only might the preferred driving style differ between people, but each individual person might also differ in their preferences depending on the situation and circumstances. Based on factors such as their emotional state and changing preferences due to experience, their preferred driving style might be completely different.

Preferred driving style includes, among others, the driving speed and rate of acceleration, and adapting an AV's driving style to each driver might be an important step in the acceptance of automated cars (Sun et al., 2020; Trende et al., 2019). The driving style that a driver would favor for an AV might not automatically be the same as when they would drive in a manual car (Bellem et al., 2018). Several studies have shown that replicating the manual driving style of the driver does not lead to a better user experience and that most users prefer a more 'defensive' driving style (see e.g. Basu et al., 2017; Ekman et al., 2019; Hasenjager & Wersing, 2017 or Yusof et al., 2016). A study by Basu et al. (2017), showed that users chose a defensive driving style that they thought matched their own, but actually did not match. This finding indicates that users are not good at objectively evaluating their driving style. A better and more personalized approach is needed that allows for 'in the moment' adaptations: e.g. adapting vehicle speed to (real-time) user trust.

In order to avoid misuse and disuse, the driver's trust should, ideally, match the reliability of the AV (Walker et al., 2019). In a recent study, Sun et al. (2020) found that the performance of an AV was considered trustworthy when it used smooth acceleration and a driving speed that participants deemed suitable to their preferred driving style. A study by Yan et al. (2017) researched the effect of adapting lane-changing assistance systems to user uncertainty on trust. They found that participants' trust levels increased when the lane-changing system adapted to their uncertainty state. Therefore, the aim should be to develop AVs which are calibrated to each person, in order to establish appropriate trust.

Current study

The goal of the present study is to research the effect that personalizing speed has on user trust in AVs. Thus, the research question of this study is: Does adapting a vehicle's speed to the user's trust lead to higher trust levels in individuals that tend to have low trust in automation? We expected that, when compared to two control groups where the *speed of the AV stays* *constant* and where the *speed of the AV randomly changes*, adapting the vehicle's speed to the current user's trust would lead to higher trust levels in individuals that tend to distrust automation. Additionally, we expected the experimental group, for which the *speed of the AV adapts to real-time user trust*, to show lower physiological arousal than the two control groups.

Methods

Participants

Participants were recruited utilizing the convenience sampling method. Most of the participants were students from the University of Twente and were recruited via the internal test subject pool of the University. Upon completion of the study, participants received a total of 2 credit points as a reward for their participation. In addition to the students from the University of Twente, the researcher recruited additional participants from his own network. A total of 50 participants completed the study, however, 5 participants were excluded from the data set due to experiencing motion sickness or because of inaccurate data. The final data set was made up of a total of 45 participants with 15 males and 30 females. They ranged from 18 to 25 years old (M = 20.6, SD = 1.9) and most were from Germany (n = 27) and the Netherlands (n = 12). Other participants' nationalities included Romanian (n = 1), Bulgarian (n = 1), Spanish (n = 1), Italian (n = 1), Hungarian (n = 1) and Slovenian (n = 1). Participants' driving experience varied from 0.5 to 8.5 years (M = 3.05, SD = 2.03) and they indicated that, on average, they drove every day (n = 9), twice per week (n = 11), once per week (n = 9), once per month (n = 14) and never (n = 2).

Each participant was randomly assigned to one of three groups, with each group being composed of 15 participants. To be eligible for participation in the study, participants had to fulfill the following criteria: sufficient understanding of English, a valid driver's license, have normal or corrected to normal vision and no colorblindness, should not easily suffer from motion sickness, and should have never driven or have been a passenger in an automated vehicle. Most importantly, participants needed to have low initial trust in AVs (this was measured by a pre-questionnaire before participants were invited to participate). Low trusting participants were recruited because, compared to individuals with higher levels of initial trust, larger changes were expected in participants with lower levels of initial trust.

Apparatus and Materials

Pre- and Post-test questionnaire

This research utilized a pre- and post-test questionnaire to measure participants' initial trust level toward automated vehicles. The pre- and post-test questionnaires were a modified version of the Empirically Derived Trust Scale (Jian et al., 2000) and consisted of seven questions on which participants could indicate their answers on a 7-point Likert scale (see Walker et al., 2018). Additionally, the pre-test questionnaire also asked for general demographic questions such as age, gender, nationality, years of driving experience, and handedness. The post-test questionnaire asked almost the same questions as the pre-test questionnaire, only that the trust questions were worded in the past tense to ensure that participants were answering the questions with the experience of having just driven in an automated vehicle in mind. Furthermore, they were asked whether they had observed any speed changes during the experiment. See Appendix A for the full questionnaires.

Consent form

Before participants were able to fill out the pre-questionnaire, they were asked to give their written consent to participate in the experiment. Participants needed to confirm that they did not suffer from any neurologic, psychiatric, or psychological condition, that they did not have any color vision deficits, that they understood that they could terminate their participation immediately if they felt uncomfortable or unwell, and that they were participating in a sober state and of their own volition. See Appendix B for the consent form.

Automated vehicle simulator

This research used a fixed base simulator (see Figure 2). As participants drove in a Level 4 automated vehicle, which did not require any driver input, they used neither the steering wheel nor the pedals. The simulation was programmed in Unity by the BMS lab at the University of Twente. In the simulation, the participants sat in a regular 5-seater passenger car that drove them around a test track without any other traffic, pedestrians or obstacles (see Figure 3). The test track was set in a mountainous coastal area and featured straight sections, as well as curves and changes in elevation. To create an immersive and authentic driving experience, participants wore an Oculus Rift Virtual Reality headset during the experiment. This allowed participants to look around in the virtual environment. The whole program ran on a computer that, next to simulating the driving experience, also recorded the vehicle's speed as well as the number of trust changes that occurred.

Figure 2

The fixed base driving simulator used in the research

Figure 3

Participant view (left) and test track that the participants were being driven on (right)

Trust slider

The trust slider (Figure 4) was originally developed for an experiment by Walker et al. (2019) to measure feelings of safety in pedestrians when wanting to cross a road when a vehicle approached. In this study, it was used by participants to indicate their in-the-moment trust on a continuous scale on which 0 indicates no trust in the vehicle and its abilities and 100 indicates full trust in the vehicle and its abilities. A small metal rod was added to the slider casing at the halfway mark, which served as a reference point for participants that they could locate without looking. This was done because in a pilot study it was found that when wearing a VR headset participants lost track after some time of how far they could slide and on which position the slider was placed exactly. The trust slider was connected to the computer and provided, at about 6 Hz, the in-the-moment trust values that were needed to induce speed changes.

Figure 4

The continuous slider developed by Walker et al. (2019) used to indicate the trust in the automated vehicle and its abilities

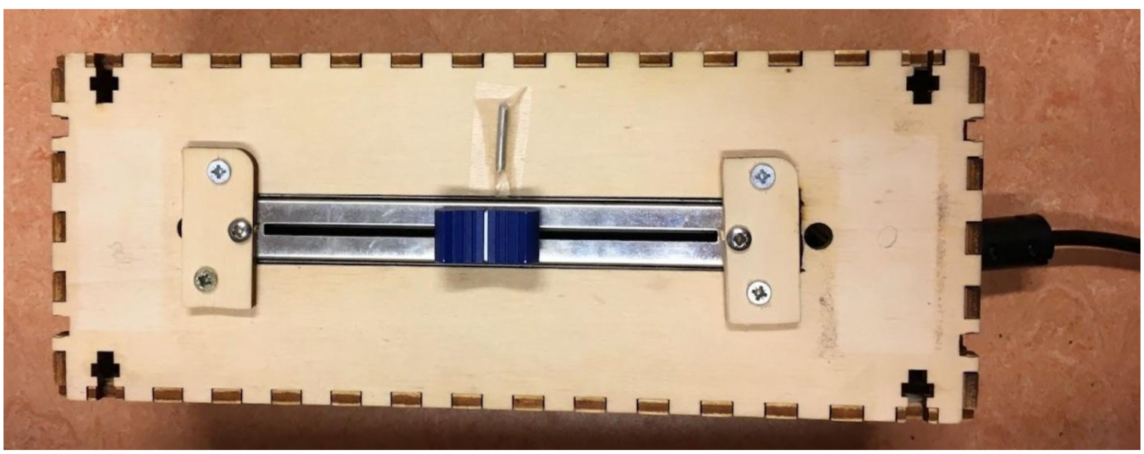

Empatica E4

The Empatica E4 is a wristband that measures several types of physiological data. In the present experiment, it was used to record the participant's skin conductance to make inferences about the participant's stress level and thereby their trust. It was put on the participant's non-dominant hand and recorded electrodermal activity at a sampling rate of 4 Hz.

Design

The study employed a between-subjects design with three different experiment groups described below serving as the independent variable. The dependent variables used in this experiment were in-the-moment trust as measured by the trust slider, difference in trust score as measured by the pre- and post-test questionnaires, and skin conductance as measured by the EDA wristband.

Task

Participants were asked to ride around the test track in the automated vehicle for 15 minutes while indicating their in-the-moment trust level throughout the drive. This research involved three conditions, one experimental condition and two control conditions, with each being carried out by one of the three participant groups. In the experimental condition, hereafter referred to as the *adaptive speed condition*, the speed of the car adapted to the real-time trust (as measured by the trust slider) of the participant. In other words, whenever there was a significant change in trust, the speed of the car was increased or decreased by 5 km/h. A change in trust was considered significant when the trust value of the participant, as measured by the trust slider, increased or decreased by 10% within a 3-second time period. Once the program determined a change in trust to be significant, the speed of the car increased or decreased with a five-second delay to make the change in speed less obvious to participants. After the five-second delay, the new speed was reached immediately. For a more detailed explanation of the variable settings such as a more comprehensive technical description of the different variables used in this research set-up, see Kowalski (in preparation). In the two control conditions, the in-the-moment trust as measured by the trust slider did not affect the speed of the vehicle, however, both speed and trust were still recorded. In the first control condition, referred to as the *fixed speed condition*, the speed of the vehicle did not change throughout the whole drive and the vehicle speed was set at a constant speed of 20 km/h in the simulator (please note that the simulated speed was not calibrated and a simulated speed of 20 km/h was probably perceived as much faster by the participant). In the second control condition, referred to as the *semi-random speed condition*, the speed changes were manually induced once every minute by the researcher. The induced changes were based on a pre-determined semi-random list of speed changes (see Appendix C for the randomized lists and the code used to create the lists). There were a total of

15 speed changes that occurred in the *semi-random control condition* while there were an average of 50 speed changes in the *adaptive speed condition* (i.e. experimental group). Participants in all three groups did not see the vehicle speed on the speedometer so as to not make them aware of any speed changes happening.

Procedure

Before people were invited to the driving simulator study, they were asked to fill in a pre-questionnaire to determine their initial trust level in automated vehicles. They were asked to rate their trust on a 7-point Likert scale by answering seven questions related to automated vehicle trust. If their trust level was too high (i.e. average >4), participants were thanked for their participation and their data were deleted. If their trust level was sufficiently low (i.e. average <=4) they were invited to participate in the actual driving simulator study in a simulator room at the University of Twente. Participants were, however, not informed about their trust level before the drive.

After participants were welcomed to the lab, they were asked to fill in a COVID-19 questionnaire to make sure they were not experiencing any symptoms, had not contracted COVID-19 within the past 14 days, and had disinfected their hands beforehand. Next, they sat down in the driving simulator and were briefed on the research procedure. Participants were told that they had to ride in the automated vehicle for 15 minutes while indicating their in-the-moment trust levels via the trust slider, whose usage was explained to them in detail. Additionally, they were also explicitly informed that there would be no takeover request from the car and that they could stop the study at any point without explanation. While they were told about the procedure a wristband was attached to the participant. After these explanations, participants put on the VR headset and adjusted it on their heads so they could see everything sharply and it would fit comfortably. They were then handed the trust slider and were positioned in the driver's seat of the virtual vehicle in a way that most closely resembled their position in a real car. Once everything was set up, participants began riding in the automated car for 15 minutes while always indicating their in-the-moment trust levels with the trust slider.

After the 15 minute drive in the simulator, participants were asked to fill in the post-questionnaire to indicate their trust in the automated vehicle that they had just driven in. Once they had finished the experiment, they were debriefed by explaining the exact purpose of the research to them as well as which experiment group they had been in. As the last step, participants were asked questions regarding their driving experience. They were asked about their nausea levels throughout the drive, when on the track they felt the most distrustful of the vehicle and its abilities as well as when they felt the most trustful, and if they had any other comments regarding the driving experience and how it could be improved. Next, they were thanked for their participation and were granted their two credits as a reward. In total, the whole study took participants about 30 minutes to complete.

Results

Changes in trust-score

Before the data were analyzed, the difference in overall trust between the pre- and post-questionnaire was calculated for each participant. For the trust score difference, first, the scores of the seven items of the pre-questionnaire (with items one and five reverse-scored) were averaged, which led to the pre-questionnaire trust score. The post-questionnaire trust score was calculated in the same manner. The change in trust score was calculated by subtracting the trust score of the pre-questionnaire from the trust score of the post-questionnaire.

To assess the effect of driving conditions on changes in trust score as well as test for changes in trust in each individual group, a 3 (conditions) x 2 (pre- and post-trust scores) mixed factors ANOVA was conducted. Before running the ANOVA analysis, the assumptions of sphericity, as well as homogeneity of variance for the trust scores, were checked. Mauchly's Test of Sphericity was used to assess the normality of the data. As there were only two levels of repeated measures, sphericity can be assumed under these conditions. Homogeneity of variance was assessed by using Levene's Test, which showed that equal variance could be assumed for the post-test scores ($p = .44$). For the pre-test scores, homogeneity of variance could not be assumed ($p = .01$). However, according to Hsu (1996), even without fulfilling the requirement of homogeneity of variance we can still interpret the post hoc test results as those are independent of the requirements of the ANOVA.

The mixed factors ANOVA indicates that there was a significant increase between the pre- and post-trust scores, $F(1, 42) = 66.7$, $p < .001$, $\eta_p^2 = .61$. Bonferroni corrected post hoc tests revealed a significant difference between the pre- and post-trust scores in all three groups (all *ps* < .001). See Figure 5 for pre- and post-trust scores between the three conditions. There was, however, no significant interaction between Conditions and Pre-/Post-trust scores, $F(2, 42) = .84$, $p = .44$, $\eta_p^2 = .04$. In other words, this result indicates that, between the three conditions, there was no difference in trust score change after having experienced an AV.

Figure 5

Average Trust Score for the Pre- and Post-trust score as Plotted by the Three Different Experimental Conditions. Error Bars Represent Two Standard Errors.

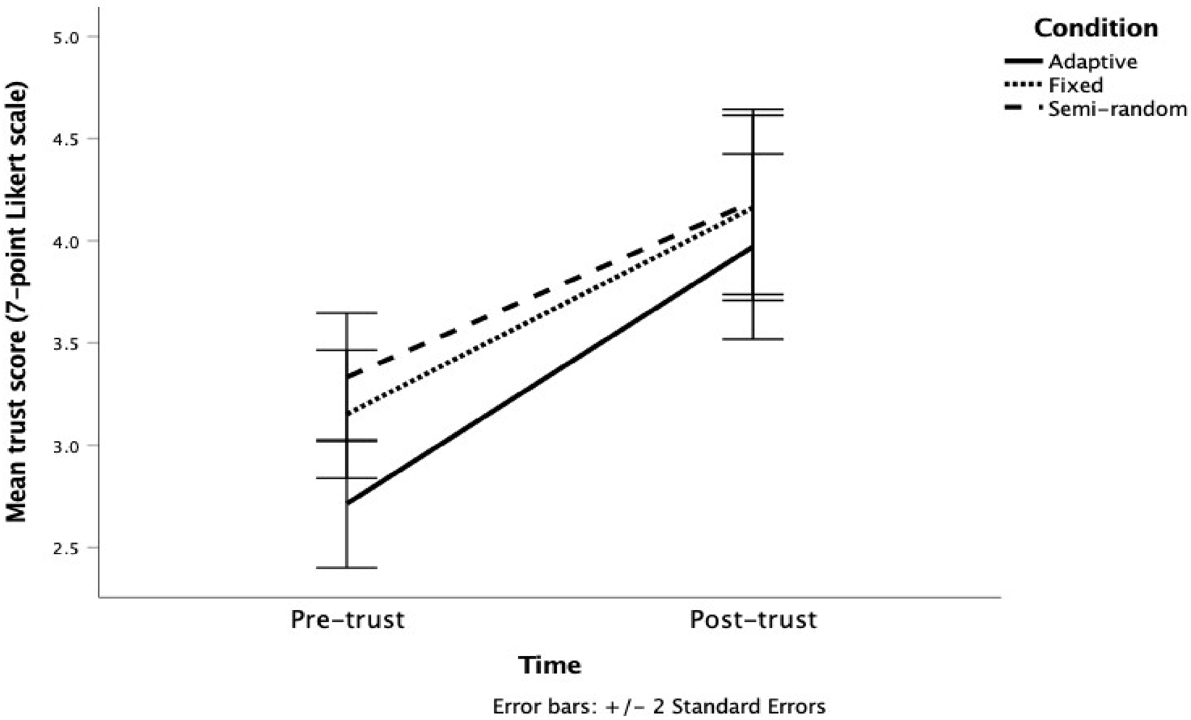

Additionally, a one-way ANOVA was conducted to assess the effect of driving conditions on post-trust scores. Before running the one-way ANOVA analysis, the assumptions of normality, as well as homogeneity of variance for the trust-score change data, were checked. The Shapiro-Wilk test ($\alpha > .05$) was used to assess the normality of the data and showed that the trust score difference was normally distributed for all three conditions. Homogeneity of variance was assessed by using Levene's Test, which showed that equal variance could be assumed ($p = .44$). Results showed that the post-trust scores were not significantly different between the different driving conditions, $F(2, 42) = .28$, $p = .76$, $\eta^2 = .01$ (see Figure 5).

Trust slider value

The average trust score, measured as trust slider position, was calculated for each participant by averaging all trust slider measurements that were collected across a period of 15 minutes. Before conducting a one-way ANOVA analysis, the trust slider values were checked on the assumptions of normality as well as homogeneity of variance. The Shapiro-Wilk test ($\alpha > .05$) was used to assess the normality of the data and showed that the trust score difference was normally distributed for all three conditions. Homogeneity of variance was assessed by using Levene's Test, which showed that equal variance could be assumed ($p = .31$).

A one-way ANOVA was conducted to assess the effect of driving conditions on overall trust slider values. Participants in the *adaptive speed condition* had an average trust slider value of 62.73% (*SD* = 14.44), participants in the *fixed speed condition* an average of 65.53% (*SD* = 9.07), and participants in the *semi-random speed condition* an average of 60.07% (*SD* = 14.74). Results showed that the trust slider scores were not significantly different between the different driving speed conditions, $F(2, 42) = .66$, $p = .52$, $\eta^2 = .03$. For an indication of how trust evolved over time instead of the average trust after 15 minutes see Figure 6.

Figure 6

Average Trust Slider Score (as Measured by Trust Slider Position) per Minute as Plotted by the Three Different Experimental Conditions.

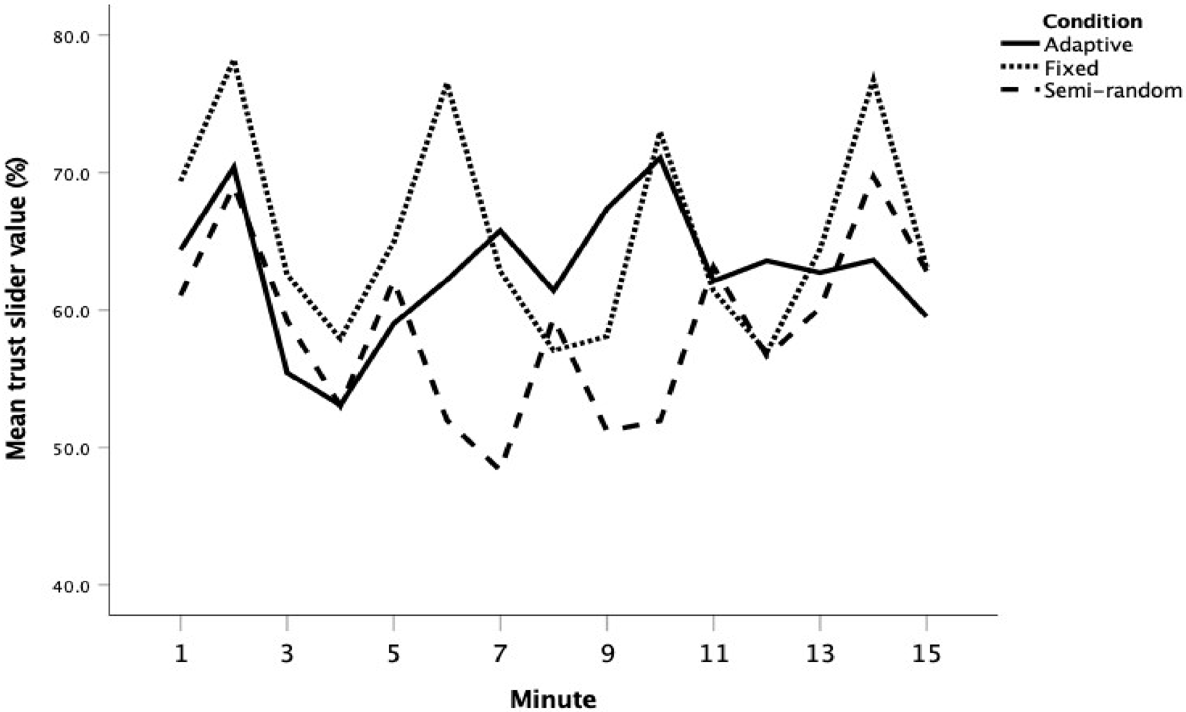

Electrodermal activity

The average electrodermal activity score, as measured by the EDA wristband, was calculated for each participant by averaging all skin conductivity measurements that were obtained from the EDA wristband. The EDA values were transformed by calculating each participant's baseline (average of the first minute of EDA measurements) and subtracting it from all EDA values.

Before running a one-way ANOVA analysis, the assumptions of normality, as well as homogeneity of variance of the EDA data, were checked. The Shapiro-Wilk test ($\alpha > .05$) was used to assess the normality of the data and showed that the trust score difference was not normally distributed for all three conditions ($p < .05$). However, one-way ANOVAs are very robust against violations of normality (see e.g. Blanca et al., 2017; Salkind, 2010; Schmider et al., 2010). Thus, it was decided that this violation will be ignored for this one-way ANOVA analysis. Homogeneity of variance was assessed by using Levene's Test, which showed that equal variance could be assumed ($p = .6$). A one-way ANOVA was conducted to assess the effect of driving conditions on overall electrodermal activity. See Figure 7 for the different EDA scores between the groups. Results showed that the EDA values were not significantly different between the different driving speed conditions, $F(2, 42) = .67$, $p = .52$, $\eta^2 = .03$.

Figure 7

Average EDA Value (Baseline Transformed) as Plotted by the Three Different Experimental Conditions. Error Bars Represent Two Standard Errors.

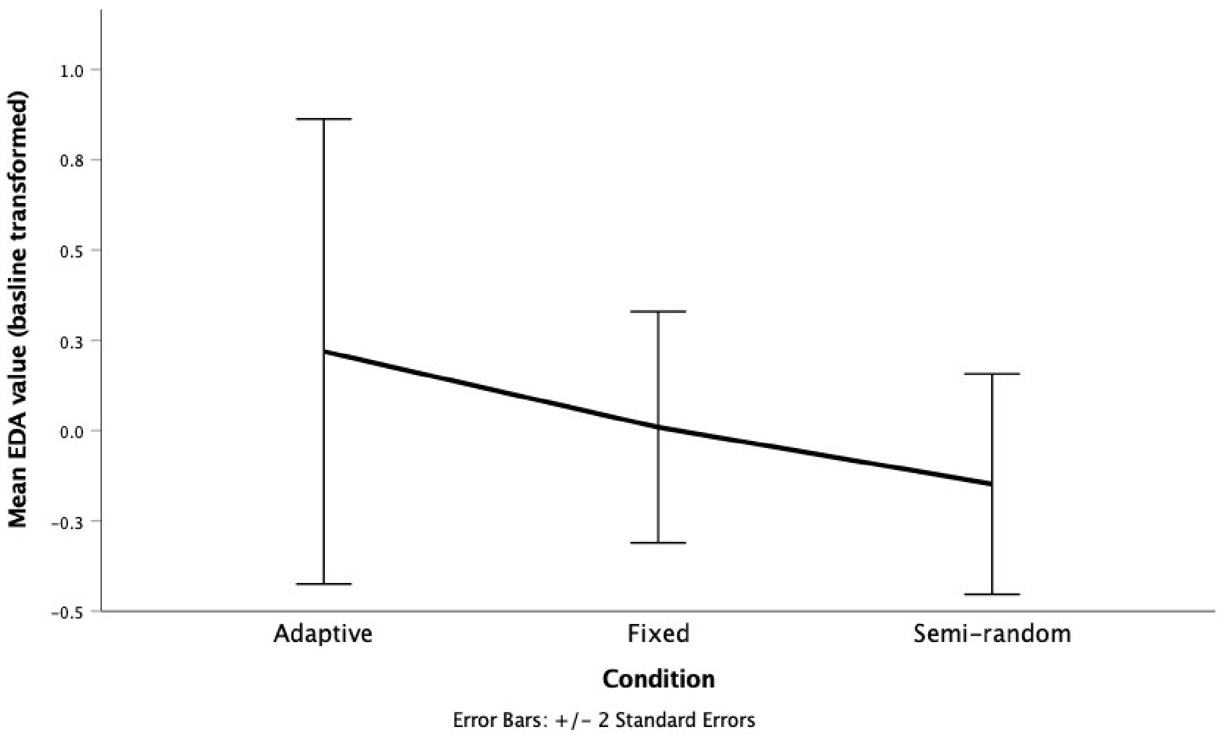

Discussion

This study aimed to investigate whether adapting an AV's speed to user trust leads to higher trust levels in individuals that tend to have low trust in automation. The following two outcomes were hypothesized. Firstly, it was expected that when compared to the two groups in which the speed of the AV stays constant and where the speed randomly changes, adapting the AV's speed to the user's trust would lead to higher trust levels in individuals who tend to distrust automation. Secondly, it was expected that the adaptive speed group, in which the speed of the

AV adapts to user trust, would show lower physiological arousal compared to the two control groups. To test the two hypotheses, participants in this study were randomly assigned to one of three groups and were asked to take a 15 minute ride in a simulated Level 4 AV. Both their initial trust level before the drive, as well as their trust level after taking the ride, were measured using a questionnaire. Furthermore, during the ride, their EDA was measured and they continuously rated their level of trust in the AV and its capabilities. This study is, to the best of our knowledge, the first to investigate adapting an AV's speed to user trust, however, we were unable to find results that support the aforementioned hypotheses.

Trust levels

The results confirmed that a difference between the pre- and post-test scores was found for all groups. This indicates that no matter which group a participant belonged to, the trust level of participants increased after having experienced a ride in an AV. A possible explanation could be that participants were specifically selected for having low trust in automated vehicles and that being merely exposed to an AV for the first time was enough for an increase in trust to occur. This is also in line with findings from Clement et al. (2022) and Hartwich et al. (2019) who found that participants experienced an increase in trust after their first time riding in an AV. Participants see that the AV behaves perfectly without making mistakes and that it behaves more reliable than they may have expected. This might have been enough for an increase in trust to occur between the pre-and post-questionnaire.

However, while there was a difference between pre-and post-test trust scores for all groups, there was no difference between the three groups in overall trust score change after having experienced an AV. The results did not confirm the hypothesis that adapting the vehicle speed to user trust leads to higher trust levels in individuals that tend to distrust automation. Neither the values obtained from the trust slider nor the trust score difference calculated from the pre- and post-trust questionnaires seemed to provide the evidence supporting the expected outcome. This suggests that there were no higher trust levels in individuals for whom the vehicle speed was adapted to user trust compared to the two groups where the speed did not adapt. A possible explanation might be that people perceived the driving style as too uncomfortable. In other words, participants did not trust the AV more in the group where the speed adapts to user trust because the speed of the car was too high in certain situations. Almost all participants mentioned that they felt the most unsafe when going through a (steep) curve as well as encountered sections where they could not see what was coming up, e.g. during a steep incline, since they deemed the speed that they were going at to be too high. This is in line with Shahrdar et al. (2019), who stated that an aggressive driving style results in reduced trust. Driving too fast might be seen as an instance of aggressive driving. Furthermore, in the present study 11 out of 15 participants in the fixed speed condition reported that they observed speed changes - even though there were none. After following up with a question on where they thought those speed changes had occurred they mentioned that they felt like the car was speeding up in curves or during a steep decline. Since people reported mostly trusting the AV on the straight parts of the track, their lowered trust when encountering a steep decline or a curve suggests that people would usually reduce their speed in those situations. This might be due to the fact that in non-automated driving people would press the brakes and slow down in situations such as approaching a curve (Vos et al., 2021). As this is not the case in the AV, people perceived it as going too fast and they felt unsafe. They, therefore, had low trust when encountering those parts of the track.

Electrodermal activity

The second expected outcome of having lower physiological arousal in the group where the speed of the AV adapts to the user's trust was also not confirmed by the results. The experiment yielded no significant differences in overall EDA between the three groups, which indicates that no one group had higher or lower physiological arousal than the other and that adapting speed to user trust does not lead to lower physiological arousal. The previously mentioned explanation of the inappropriate AV driving style might also be the explanation for this outcome. Participants might have deemed the driving style of the AV not defensive and safe enough. As drivers prefer defensive driving styles (Basu et al., 2017; Ekman et al., 2019; Hasenjager & Wersing, 2017; Yusof et al., 2016), too high speeds in situations where the participant felt like a slower speed would be more appropriate, might have led to no difference in EDA between the three groups. An additional explanation might have been that the drive was simulated. In other words, participants understood that there was danger, but since they were not expecting any real consequences in the simulated environment, there was no/too little of an emotional or physical response.

Limitations and future research

This study suffered from a number of limitations that are important to be discussed. Firstly, there are some limitations surrounding how the trust slider works and is used by participants. If a participant had the trust slider at the extremes of 100% or 0% trust for a long period of time, the speed did not further increase or decrease. This might have led to distorted results of participants having a higher or lower level of trust than their real trust. It is especially important to decrease the vehicle speed when the participant has the trust slider set to 0% for an extended time since the participant should not feel unsafe while riding an AV. Therefore, in future studies, the speed should further increase or decrease when the trust slider is set to 100% or 0% respectively for an extended period. Furthermore, due to the between-subjects design of the study, inter-individual differences cannot be ruled out. For example, some participants did not the whole spectrum and only used the bottom third to indicate their trust, compared to others who used the whole spectrum of the trust slider. Those who used only a small spectrum might have experienced much more speed changes than the participants that used the whole slider spectrum. Thus, a future study should account for inter-individual differences by for example using a repeated measure within-subjects design and perhaps more standardized instructions on how to use the slider.

Another limitation is related to the EDA measurements since the measurements might have been influenced by people feeling slight motion sickness. Even though a prerequisite of this study was that participants should not easily feel motion sick, as well as participants being excluded from the dataset for experiencing too much nausea, some participants said that they felt slight nausea at certain points during the ride, which might have increased their EDA. The EDA might even have been influenced by factors such as participants coming in on a warmer day on their bike and therefore were starting to sweat more during the experiment. Thus, the EDA measurements could be complemented by another physiological measurement such as gaze behavior, as this can give an extra unobtrusive indication of user trust in the AV. This is in line with the study of Walker et al. (2019), who suggested a combination of a physiological measurement (i.e. EDA) and gaze behavior can give a more precise indication of trust than one individually. As a lot of simulator studies use VR glasses, some of which have an in-built eye-tracking option, also measuring gaze behavior would be quite easy and would not involve adding more measuring devices to the participant. Additionally, a follow-up study should look into reducing nausea in the driving simulator by, for example, using a motion-base simulator that can simulate lateral forces. Currently, we only controlled for motion sickness by telling participants that they should point out once they experience some motion sickness, as well as by inquiring about the nausea levels after the drive. It is crucial that a standardized measure of motion sickness should be employed in the future, both during the recruiting part as well as during the study, to avoid possibly harmful consequences to participants.

Conducting research in a simulator always takes away the risks of riding in a real AV. In a simulator, there are never any harmful consequences or injuries that one could sustain compared to in a real vehicle. Knowing this might influence the participant in a way that they trust the AV more than they would in a real automated vehicle. While simulators protect subjects from harm, real-life validation is needed. Therefore, after more follow-up studies in a simulator (possibly in a motion-base simulator), the results should be validated in a real AV. Moreover, on the simulator track there was no traffic, no intersections, no traffic lights, no stops, nor did the AV perform any overtakes, it was just 'simple' continuous driving. It might be good to include at least some of those features in the future, in order to simulate a driving scenario that is closer to a real driving environment.

The last limitation of this study was that the study was not representative of a larger population. Almost all participants were students (from a technical university), who already might have had a higher trust towards automation in general as well as being more technology savvy than the population at large. Although this study tried to control for that by making low trust (as measured by the pre-questionnaire) a prerequisite, a future study might benefit from recruiting a participant sample that is more representative of the general population.

General future recommendations

After talking to and observing participants a few things stood out. As mentioned, almost all participants pointed out that they felt the most unsafe when going through a (steep) curve as well as when they encountered sections where they could not see what was coming up, e.g. during a steep incline. Thus, generally decreasing the speed by X% in sections where a lot of participants felt they were going too fast and experienced low trust might lead to higher trust levels. This would be in line with earlier reported studies where they found that participants preferred a more defensive driving style (e.g. Basu et al., 2017; Ekman et al., 2019; Hasenjager & Wersing, 2017; Yusof et al., 2016).

Currently, judgments on where on the track the most significant increases or decreases in trust occurred were only based on visual observation as well as participant questioning. To increase accuracy, a follow-up study might profit from tracking the location of the vehicle to infer where on track the most significant changes in trust took place. Additionally, this study was not able to properly analyze the trust slider values' and EDA measurements' development of trust over time. Future studies might benefit from looking into a time series analysis and analyzing potential changes and peaks paired with the location data of the vehicle.

On top of that, as mentioned above, a simulated vehicle speed of 20 km/h did not match a real speed of 20 km/h, which could have led to bias for some participants. At the moment, we cannot accurately say how fast the speed was actually perceived by participants. To avoid biasing participants a future study should therefore make sure that the simulated speed is calibrated and also matches the real vehicle speed. Moreover, in the current study speed changes are induced in steps of 5 km/h and with a 5-second delay. A future study might benefit from having smoother acceleration/deceleration by gradually increasing/decreasing the speed over 5 seconds with 1 km/h added/taken each second. This would hopefully make speed changes even less obvious, as well as start adding/reducing speed right when a speed change is needed.

Even though the results do not seem to support the effect of adapting the vehicle speed of an AV to the user's trust resulting in higher trust levels compared to groups where the speed did not adapt to user trust, future studies might still be able to confirm this effect. Studies such as Sun et al. (2020) and Yan et al. (2017) might not have researched adaptive speed changes specifically, but they have shown, that if AV systems are adaptive and personalized, they can increase trust. The adaptive speed change settings in this study likely need some more fine-tuning for the effect of increased trust due to adaptive speed to appear. Thus, incorporating the above-mentioned recommendations might provide the necessary information for a follow-up study that further investigates the effect of adaptive speed changes.

Conclusion

With more and more research and development going into the domain of automated vehicles, it becomes increasingly important to investigate good ways to facilitate adequate trust in AVs and their capabilities. This research contributes to the field with a preliminary investigation of whether adapting a vehicle's speed to the user's trust leads to higher trust levels in individuals that tend to have low trust in automated vehicles. While no evidence supporting this statement was found within the parameters as designed, this study attempted to provide insights into a topic that, up until this point, has received little to no research. The findings, as well as recommendations in this study, should be used in facilitating follow-up studies and their design.

References

Akash, K., Hu, W.-L., Jain, N., & Reid, T. (2018). A classification model for sensing human trust in machines using EEG and GSR. *ACM Transactions on Interactive Intelligent Systems*, *8*(4), 1–20. https://doi.org/10.1145/3132743

Azevedo-Sa, H., Jayaraman, S. K., Esterwood, C. T., Yang, X. J., Robert, L. P., Jr., & Tilbury, D. M. (2020). Real-Time estimation of drivers' trust in automated driving systems. *International Journal of Social Robotics*, *13*(8), 1911–1927. https://doi.org/10.1007/s12369-020-00694-1

Basu, C., Yang, Q., Hungerman, D., Singhal, M., & Dragan, A. D. (2017, March 6). Do you want your autonomous car to drive like you? *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. http://dx.doi.org/10.1145/2909824.3020250

Bellem, H., Thiel, B., Schrauf, M., & Krems, J. F. (2018). Comfort in automated driving: An analysis of preferences for different automated driving styles and their dependence on personality traits. *Transportation Research Part F: Traffic Psychology and Behaviour*, *55*, 90–100. https://doi.org/10.1016/j.trf.2018.02.036

Blanca, M. J., Alarcón, R., Arnau, J., Bono, R., & Bendayan, R. (2017). Non-normal data: Is ANOVA still a valid option? *Psicothema*, *29*(4), 552–557. https://doi.org/10.7334/psicothema2016.383

Clement, P., Veledar, O., Könczöl, C., Danzinger, H., Posch, M., Eichberger, A., & Macher, G. (2022). Enhancing Acceptance and Trust in Automated Driving through Virtual Experience on a Driving Simulator. *Energies*, *15*(3), 781. https://doi.org/10.3390/en15030781

Ekman, F., Johansson, M., Bligård, L.-O., Karlsson, M., & Strömberg, H. (2019). Exploring automated vehicle driving styles as a source of trust information. *Transportation Research Part F: Traffic Psychology and Behaviour*, *65*, 268–279. https://doi.org/10.1016/j.trf.2019.07.026

Ekman, F., Johansson, M., & Sochor, J. (2018). Creating appropriate trust in automated vehicle systems: A framework for HMI design. *IEEE Transactions on Human-Machine Systems*, *48*(1), 95–101. https://doi.org/10.1109/thms.2017.2776209

Hartwich, F., Witzlack, C., Beggiato, M., & Krems, J. F. (2019). The first impression counts – A combined driving simulator and test track study on the development of trust and acceptance of highly automated driving. *Transportation Research Part F: Traffic Psychology and Behaviour*, *65*, 522–535. https://doi.org/10.1016/j.trf.2018.05.012

Hasenjager, M., & Wersing, H. (2017, October). Personalization in advanced driver assistance systems and autonomous vehicles: A review. *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*. http://dx.doi.org/10.1109/itsc.2017.8317803

Hoff, K. A., & Bashir, M. (2014). Trust in automation. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *57*(3), 407–434. https://doi.org/10.1177/0018720814547570

Hsu, J. (1996). *Multiple comparisons: Theory and methods*. Chapman & Hall.

Jian, J.-Y., Bisantz, A. M., & Drury, C. G. (2000). Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics*, *4*(1), 53–71. https://doi.org/10.1207/s15327566ijce0401_04

Kaur, K., & Rampersad, G. (2018). Trust in driverless cars: Investigating key factors influencing the adoption of driverless cars. *Journal of Engineering and Technology Management*, *48*, 87–96. https://doi.org/10.1016/j.jengtecman.2018.04.006

Kaye, S.-A., Somoray, K., Rodwell, D., & Lewis, I. (2021). Users' acceptance of private automated vehicles: A systematic review and meta-analysis. *Journal of Safety Research*, *79*, 352–367. https://doi.org/10.1016/j.jsr.2021.10.002

Khawaji, A., Zhou, J., Chen, F., & Marcus, N. (2015, April 18). Using galvanic skin response (GSR) to measure trust and cognitive load in the text-chat environment. *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. http://dx.doi.org/10.1145/2702613.2732766

Kowalski, M. M. (in preparation). *Automated vehicle behaviour: a driving simulator study*. [Unpublished master thesis]. Faculty of Cognitive Psychology and Ergonomics, University of Twente.

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *46*(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

Morris, D. M., Erno, J. M., & Pilcher, J. J. (2017). Electrodermal response and automation trust during simulated self-driving car use. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *61*(1), 1759–1762. https://doi.org/10.1177/1541931213601921

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *39*(2), 230–253. https://doi.org/10.1518/001872097778543886

Pettersson, I., & Karlsson, I. C. M. (2015). Setting the stage for autonomous cars: A pilot study of future autonomous driving experiences. *IET Intelligent Transport Systems*, *9*(7), 694–701. https://doi.org/10.1049/iet-its.2014.0168

Reyes, A. (2021, December 13). Mercedes-Benz wins world's first approval for Level 3 autonomous cars: What's that mean? *SlashGear*. https://www.slashgear.com/mercedes-benz-wins-worlds-first-approval-for-level-3-autonomous-cars-whats-that-mean-13702207/

SAE. (2021a). *SAE J3016 Levels of Driving Automation*. https://www.sae.org/binaries/content/assets/cm/content/blog/sae-j3016-visual-chart_5.3.21.pdf

SAE. (2021b). *(R) Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. 1–41. https://doi.org/10.4271/J3016_202104

Schmider, E., Ziegler, M., Danay, E., Beyer, L., & Bühner, M. (2010). Is it really robust? *Methodology*, *6*(4), 147–151. https://doi.org/10.1027/1614-2241/a000016

Shahrdar, S., Park, C., & Nojoumian, M. (2019, January 27). Human trust measurement using an immersive virtual reality autonomous vehicle simulator. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. http://dx.doi.org/10.1145/3306618.3314264

Song, J., Colasante, T., & Malti, T. (2019). Taming anger and trusting others: Roles of skin conductance, anger regulation, and trust in children's aggression. *British Journal of Developmental Psychology*, *38*(1), 42–58. https://doi.org/10.1111/bjdp.12304

Sun, X., Li, J., Tang, P., Zhou, S., Peng, X., Li, H. N., & Wang, Q. (2020). Exploring personalised autonomous vehicles to influence user trust. *Cognitive Computation*, *12*(6), 1170–1186. https://doi.org/10.1007/s12559-020-09757-x

Trende, A., Gräfing, D., & Weber, L. (2019, September 21). Personalized user profiles for autonomous vehicles. *Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications: Adjunct Proceedings*. http://dx.doi.org/10.1145/3349263.3351315

United Nations General Assembly. (2020). *Improving global road safety* (pp. 1–9). https://undocs.org/en/A/RES/74/299

Vos, J., Farah, H., & Hagenzieker, M. (2021). Speed behaviour upon approaching freeway curves. *Accident Analysis & Prevention*, *159*, 106276. https://doi.org/10.1016/j.aap.2021.106276

Walker, F. (2021). *To trust or not to trust? : Assessment and calibration of driver trust in automated vehicles* [Doctoral dissertation, University of Twente]. http://dx.doi.org/10.3990/1.9789055842766

Walker, F., Steinke, J., Martens, M. H., & Verwey, W. B. (2020). Do engineer perceptions about automated vehicles match user trust? Consequences for design. *Transportation Research Interdisciplinary Perspectives*, *8*, 100251. https://doi.org/10.1016/j.trip.2020.100251

Walker, F., Wang, J., Martens, M. H., & Verwey, W. B. (2019). Gaze behaviour and electrodermal activity: Objective measures of drivers' trust in automated vehicles. *Transportation Research Part F: Traffic Psychology and Behaviour*, *64*, 401–412. https://doi.org/10.1016/j.trf.2019.05.021

Yan, F., Eilers, M., Ludtke, A., & Baumann, M. (2017, June). Building driver's trust in lane change assistance systems by adapting to driver's uncertainty states. *2017 IEEE Intelligent Vehicles Symposium (IV)*. http://dx.doi.org/10.1109/ivs.2017.7995772

Yusof, N. Md., Karjanto, J., Terken, J., Delbressine, F., Hassan, M. Z., & Rauterberg, M. (2016, October 24). The exploration of autonomous vehicle driving styles. *Proceedings of the*

8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications. http://dx.doi.org/10.1145/3003715.3005455

Zhang, T., Tao, D., Qu, X., Zhang, X., Lin, R., & Zhang, W. (2019). The roles of initial trust and perceived risk in public's acceptance of automated vehicles. *Transportation Research Part C: Emerging Technologies*, *98*, 207–220. https://doi.org/10.1016/j.trc.2018.11.018

Appendix

Appendix A

Pre-test questionnaire:

Part 1:

Name:

Age:

Gender:

Nationality:

Years of driving experience (i.e. Years that passed since when you first got your driving license):

On average, how often do you drive on European roads (including Dutch roads)?

- Never
- Once per month
- Once per week
- Twice per week
- Every day

Part 2:

Through this brief questionnaire, we would like to measure your attitudes toward self-driving cars.

We understand that your knowledge of self-driving cars might be limited, so please answer based on your ideas and expectations.
Please respond as truthfully as possible, and keep in mind that there is no “correct” answer.

Please indicate your answer.

1 = not at all

7 = extremely

1. **I am cautious about self-driving cars**

1 – 2 – 3 – 4 – 5 – 6 – 7

2. **Self-driving cars are reliable**

1 – 2 – 3 – 4 – 5 – 6 – 7

3. **I would entrust my car to self-driving functions for lane keeping, lane changing, alerts following object recognition, etc.**

1 – 2 – 3 – 4 – 5 – 6 – 7

4. **I can count on self-driving cars**

1 – 2 – 3 – 4 – 5 – 6 – 7

5. **Self-driving cars can have harmful consequences**

1 – 2 – 3 – 4 – 5 – 6 – 7

6. **I trust self-driving cars**

1 – 2 – 3 – 4 – 5 – 6 – 7

7. **I assume that self-driving cars will work properly**

1 – 2 – 3 – 4 – 5 – 6 – 7

Post-test questionnaire:

Participant number:

Answer the questionnaire keeping in mind the behavior of the car you have been riding during this experiment.
Please respond as truthfully as possible, and keep in mind that there is no “correct” answer.

Your privacy is protected according to Dutch law.

Please indicate your answer.

1 = not at all

7 = extremely

1. **I was cautious about the self-driving car**

1 – 2 – 3 – 4 – 5 – 6 – 7

2. **The self-driving car was reliable**

1 – 2 – 3 – 4 – 5 – 6 – 7

3. **I would entrust my car to the tested self-driving functions (for example, lane-keeping)**

1 – 2 – 3 – 4 – 5 – 6 – 7

4. **I could count on the self-driving car**

1 – 2 – 3 – 4 – 5 – 6 – 7

5. **This self-driving car can have harmful consequences**

1 – 2 – 3 – 4 – 5 – 6 – 7

6. **I trusted the self-driving car**

1 – 2 – 3 – 4 – 5 – 6 – 7

7. **The self-driving car worked properly**

1 – 2 – 3 – 4 – 5 – 6 – 7

Appendix B

Informed consent form:

I confirm that:

- I do not suffer from any neurologic, psychiatric, or psychological condition.
- I do not have any color vision deficits and I do not wear glasses (does not include contact lenses).
- I do not suffer from motion sickness.
- I have never driven or been a passenger in an automated vehicle before (for example, Tesla with Autopilot)
- I have a driver's license.
- I understand that I can terminate my participation immediately if I feel uncomfortable or unwell.
- I participate in a sober state and of my own free will.

Appendix C

Semi-random speed changes that were induced:

Minute	Speed person 1-5	Speed person 6-10	Speed person 11-15
0	20	20	20
1	15	25	25
2	10	20	30
3	5	25	25
4	10	30	20
5	15	35	25
6	20	30	30
7	25	25	35
8	30	30	30
9	35	35	25
10	30	30	20
11	25	25	25
12	20	20	30
13	15	25	25
14	20	30	20

Script used to generate the numbers in Python:

```
import random
result = 20
list = [20]

while len(list) < 15:
    if result == 5:
        result = result +5
        list.append(result)
    if result == 40:
        result = result-5
        list.append(result)
    else:
        result += random.choice ((5, -5))
        list.append(result)
print(list)
```