Adaptive Driving Agent

From Driving a Machine to Riding with a Friend

Claudia V. Goldman[†]
General Motors
Herzliya Pituach, Israel
claudia.goldman@gm.com

Albert Harounian
Computer Science Department
Bar Ilan University
Ramat Gan, Israel
albert.harounian@gmail.com

Ruben Mergui
General Motors
Herzliya Pituach, Israel
ruben.mergui@gm.com

Sarit Kraus
Computer Science Department
Bar Ilan University
Ramat Gan, Israel
sarit@cs.biu.ac.il

ABSTRACT

The successful integration of automation in systems that affect human experiences requires the user acceptance of those automated functionalities. For example, the human comfort felt during a ride is affected by the automated control behavior of the vehicle. The challenge presented in this paper is how to develop an intelligent agent that learns its users' driving preferences and adjusts the vehicle control in real time, accordingly, minimizing the number of otherwise required manual interventions. This is a hard problem since users' preferences can be complex, context dependent and do not necessarily translate to the language of machines in a simple and straightforward manner. Our solution includes (1) a simulation test bed, (2) an adaptive intelligent interface and (3) an adaptive agent that learns to predict user's driving discomfort and it also learns to compute corrective actions that maximize user acceptance of automated driving. Overall, we conducted three user studies with 94 subjects in simulated driving scenarios. Our results show that our intelligent agent learned to successfully predict how to adjust the automated driving style to increase user' acceptance by decreasing the number of user manual interventions.

CCS CONCEPTS

• Computing Methodologies Artificial Intelligence Intelligent Agents • Machine Learning Applied Computing Driving Control • Human Computer Interaction

[†]Corresponding Author

HAI '20, November 10–13, 2020, Virtual Event, NSW, Australia

ACM ISBN 978-1-4503-8054-6/20/11…$15.00
https://doi.org/10.1145/3406499.3415067

KEYWORDS

Intelligent Agents, Adaptive Behavior, User Modeling

ACM Reference format:

Claudia V. Goldman, Albert Harounian, Ruben Mergui and Sarit Kraus. 2020. Adaptive Driving Agent: From driving a machine to riding with a friend. In *Proceedings of the 8th International Conference on Human-Agent Interaction (HAI' 20)*, November 10–13, 2020, Virtual Event, Australia. ACM, NY, NY, USA. 8 pages. https://doi.org/10.1145/3406499.3415067

1 Introduction

Advances in sensing and computational technologies pave the way for automating increasing number of driving functionalities. These engineering solutions result in improved vehicle control and in freeing the human from manually controlling the vehicle at times. Nevertheless, it is essential to recognize that humans are different, thus preferring different styles of driving when facing different routes, car occupancy, and driving contexts. Usually, default, engineering-based driving styles are pre-set to control the correct and safe performance of vehicles. There is a hidden assumption that humans will all accept this style under all circumstances. However, for different people, the same driving maneuver taken with different styles may be perceived as "too aggressive" by some, whereas others would consider it to be "too cautious" or "uncomfortable". Integrating these contextual preferences in a learning agent is a challenge. Understanding users' needs and preferences is hard since these preferences are diverse, can change over time [11] and they are contextual [21,24].

This paper attempts to solve this challenge, which requires solving three problems: (1) how can an agent interact with human drivers or passengers to get and interpret their preferences, (2) how can this agent learn from these preferences and driving contexts to predict discomfort and (3) how can this agent learn to adjust online the driving style settings of the vehicle it controls to avoid predicted discomfort. The first problem is hard since people might have difficulty in expressing their needs in such complex contexts as driving. The evaluation of users' satisfaction may be also very expensive [8] (e.g., setting cameras inside the car, analyzing the passengers' video/audio streams). Therefore, it may be more efficient to rely on inputs expressed by users in natural ways enabling them an interaction similar to one they would have with a taxi driver. Namely, we would like to develop an agent that would be able to interpret requests such as "Go faster" or "You're too close" to changes in the technical configuration of the car without requiring users to give technical details how the vehicle should be achieving these requests (e.g., "use the gas paddle less smoothly" or "limit the acceleration"). Furthermore, we would like the agent to reduce the need of the human to make explicit comments by dynamically adjusting the cars' parameters based on estimation of the human's preferences. To solve the last two problems, we present a novel learning-based agent for the automatic adjustment of a car's driving style configuration by processing and intelligently reacting to natural inputs from drivers or passengers. Our agent, named the Adaptive Car Controller Agent (ACCA), was developed and tested in a state-of-the-art realistic simulation environment. Through extensive user studies, with 94 human participants, we show that the ACCA can significantly reduce user's burden of adjusting the driving style manually to achieve acceptable satisfaction levels of comfort (measured with usability questionnaires and evaluated quantitively with the number of manual interventions required through the studies to express discomfort). We have applied a similar computational approach successfully in the thermal domain [18]. We developed an intelligent (intuitive to use) interface for drivers to change the settings of their automotive climate control system, reducing the number of manual interventions required. This interface was implemented on a tablet and used by all participants through all experiments, enabling them to choose what adjustments to make to the driving style through touching buttons and sliders with intuitive meaning.

Our adaptive solution was developed in three stages, each one contributing an essential component combined eventually to build up a complete adaptive agent solution: (1) **Automatically interpretation of user intuitive needs**: we trained an intelligent agent with human data, provided through an adaptive intelligent interface, enabling users to express their driving needs without stating specific values for all related parameters. The agent was able to interpret these human inputs into actual driving control actions. (2) **Automatically learning of human-preferred driving style settings**: our agent learned 3 human-centered driving styles from data collected in the first experiment. (3) **Automatically learning and mitigation of user discomfort prediction**: our agent learned to predict driving discomfort and learned to avoid it in real time, by adjusting the driving style of the simulated car accordingly. It computed the control actions that attained the highest likelihood of being accepted by the users under similar driving situations. Our results showed that our agent adapted its control behavior successfully to its users' expectations, increasing their acceptance of the automated control and decreasing the number of manual interventions required to accept this automated behavior.

2 Related Work

The personalization of human-vehicle interactions has been studied by researchers and practitioners over the years including: the personalization of a car's climate control system [18], the adaptation of the cruise control system [20] and automatic speed control [25] to name a few. In the more complex context of adjusting a driving style, many car's settings and their interdependencies affect the user experience (e.g., the smoothness of turning the steering wheel, or of pressing the pedals (gas / brake), acceleration rate), making existing approaches unsuitable for our task. Automating human-like driving was studied to provide personal comfort [3,4]. Our approach is distinct and novel since our agent (1) *learns* a model of human preferences for driving styles, (2) *adapts* its automated behavior online accordingly, (3) is *evaluated* with real subjects riding a simulated car and (4) *intuitively interacts* with humans overcoming difficulties encountered by drivers, when expressing what they need from the car in technical terms [6]. Natural interfaces [15] are commonly more intuitive for drivers and passengers to use and understand and can assist in personalizing the interaction [13, 17]. Most relevant to our task is the work by Geng et al. [9], who provide a scenario-adaptive driving style prediction ontology. The proposed ontology presents how a car's driving style should adapt to different traffic scenarios to meet most drivers' expectations. However, the proposed ontology is limited to a "generalized prediction" as opposed to a "personalized prediction" [19]. Namely, the ontology does not provide a car with the means to adapt to each user nor does it provide a way to adapt the prediction in real-time. Other user studies in simulated driving environments analyzed driving styles and their effect on different users [1, 2]. However, these works collected data of user acceptance via questionnaires and not by interacting with an intelligent agent in real-time, as we do. Guna et al. [10] used driving data to predict driving styles given users' activities. Mazzulla et al. [16] studied the relationship between drivers' characteristics and driving styles. In all these works, the online adaptation of the driving style is missing and is presented as the next required step. In the control domain, Senouth et al. [22] developed a fuzzy rule to adaptively modulate and assist the driver and vehicle torque to keep the lane preferred by the driver. This solution was evaluated by numerical simulations and not by real interactions with subjects as we present here. Our work is novel due to the combination of human-machine interactions with machine learning to provide one solution to interpret users' preferences and then realize these into actual control actions with the aim of increasing acceptance.

3 Simulation Set Up

Our agent, user studies and evaluations were run in a Unity-based simulation environment of automated driving. Subjects playing the role of passengers sit on a physical car seat while watching three screens that simulate what they could see through the car's windows and mirrors while driving. Users were told they were going to take a taxi-like ride with a simulated automated driver in the city of San Francisco.

3.1 Simulation Scenarios

We chose an area in San Francisco city for our simulated geographical urban area. The real content of the city was transferred into graphical content from videos taken in San Francisco. Fig. 1 shows the route arbitrarily generated in the city. Using this route, we designed four driving scenarios. Each participant experienced the same route with all eleven events, four times in random order (each simulated ride took around 5 minutes). The events, distributed along the route (see Fig. 1) include common urban events such as pedestrian crossings, traffic jams, a jaywalker suddenly crossing the street or a cyclist riding next to the car, approaching a jam, driving in a jam or open road, encountering a car leaving its parking or lane, encountering hazards on road. These events were predefined to provide a rich driving context similar to real-world events. Each one of the four rides each participant experienced, differ in the settings assigned to each one of the eleven events, creating different driving contexts. We will examine the users' reactions and desired changes to the car's driving style in these scenarios.

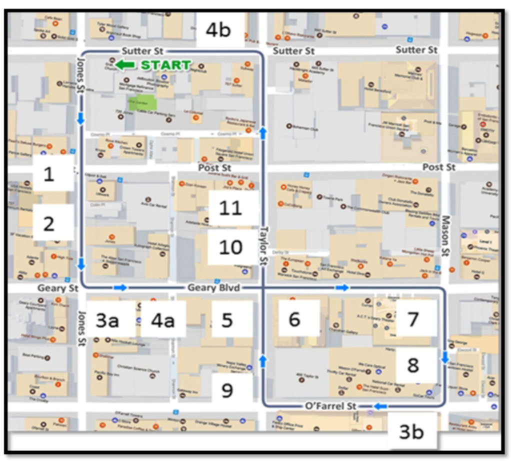

Figure 1: Simulator Route. Numbers indicate the locations of the events.

3.2 Driving Styles

A car's driving style is represented as a vector of driving parameters values. These parameters correspond to the control settings in the Unity simulator. We focused on the following ones, that affect the simulated driving behavior:

1. Gas Smooth - Determines how hard or soft to press the gas pedal. Value: [0.1, 0.9]
2. Brake Smooth - Determines how hard or soft to press the brake. Value: [0.1, 0.9]
3. Gap Distance - Minimal distance to maintain when approaching an object. Value: [3, 17]
4. Gap Time - Minimal stopping time to maintain while traveling behind another car. Value: [1, 4]
5. Forward Approach Gas Smooth - Determines how abruptly to release the gas pedal when approaching another car. Value: [0.1, 0.9]
6. Forward Approach Brake Smooth - Determines how hard or soft to press the brake pedal when approaching an object. Value: [0.1, 0.9]
7. Acceleration Rate - The acceleration applied to reach the target speed. Value: [0.1, 25]
8. Lane Centering - Determines the location of the car relative to the center of the lane. Value: [0.1, 0.9]
9. Turn Speed - Determines the desired speed when performing turns. Value: [8, 40]
10. Maximum Speed - Determines the maximum speed of the car on an unobstructed road. Value: [10, 50]

Any combination of values to these parameters would result in (slightly) different driving styles. Table 1 presents two basic driving styles, calm and active, that we defined after having tested them in the simulator. These styles will serve as the initial default driving styles in the first experiment reported below. Subjects in our experiments, see the traffic on the road and see and hear the effects of their own vehicle driving through the screens and speakers, as they would do if they were sitting in a seat next to a driver in a regular car. The goal of our agent is to learn to adjust driving styles to those preferred by humans in real time.

Parameter Name	Calm	Active
Gas Smooth	0.9	0.1
Brake Smooth	0.8	0.1
Gap Distance	10	3
Gap Time	3	1
Forward Approach Gas Smooth	0.9	0.2
Forward Approach Brake Smooth	0.8	0.2
Default Acceleration	4	25
Maximum Speed	16	50
Lane Centering	0	0
Turn Speed	13	34

Table 1: Basic (Default) Driving Styles

3.3 Human Agent Interactions

Adjusting the values of driving parameters would be easy if users could tell explicitly what they need in the language of the driving control system to attain driving comfort. However, this is not the case. In preliminary trials, participants were asked to express themselves as they would to a friend who is driving or a taxi driver. We noticed that participants were able to distinguish between the different driving styles and were able to discuss them in terms of their perceived safety and enjoyment. However, it turned out that participants were not able to satisfactorily express their preferences in terms of parameter values, resulting in many user inputs and general dissatisfaction, although they fully understand the role of each parameter. Specifically, participants were unable to determine which parameter they should change and to what extent to bring about the desired change. Moreover, we noticed that participants express different expectations based on road conditions and context (e.g., in a traffic jam vs an open road). Specifically, while it is reasonable to assume that users can distinguish between comfortable or not (e.g., feeling safe or in danger), it is unreasonable to assume that non-expert users would be able to quickly manually configure the above parameters. To that end, following these preliminary trials, we identified the following 8 terms which people often use to express their preferred driving style: 1) More Speed; 2) Less Speed; 3) More Gap; 4) Less Gap; 5) More Sport; 6) More Comfort; 7) More to Right; and 8) More to Left. It is therefore our goal in this paper to develop an automated agent that is capable of intelligently translating these natural expressions to the desired set of values to assign to the technical parameters (Section 3.2). Fig. 2 shows the interface we implemented on a Samsung tablet for getting inputs from real users during the experiments. Users could express their desired changes in driving to the experimenter and the control agent through this interface by resizing the circle, relocating the circle (4 directions) or by moving the sports slider. The users did not have to assign a value for the change requested, just the desired direction of change. Subjects would sit on a car seat and will observe 3 wide screens showing what a passenger would see from the vehicle cabin. All rides occurred in the simulator screens, so the experience was safe and did not incur any risks to the participants. Driving settings was always kept under safety constraints.

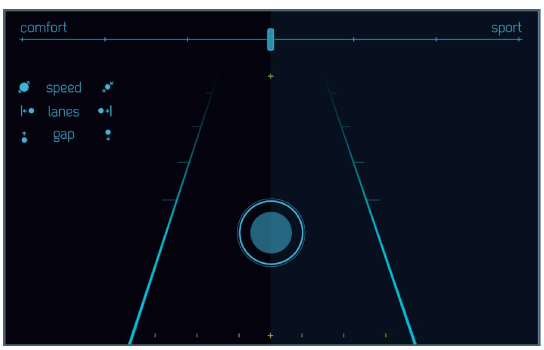

Figure 2: Adaptive Driving Agent: Natural Interface

3.4 The Adaptive Control Agent

The ACCA agent we developed was able to control the simulated vehicle through the route we defined. The agent could operate in two modes: fixed, or adaptive. The agent was always initialized with a driving style vector that determines the ranges of values of the driving parameters. The simulated vehicle applies these parameters to create the simulated physical dynamics of the driving context. In the adaptive mode, the agent could change the driving settings proactively during the ride, based on its learned models of the users' preferences. Moreover, the agent could interact with the human subject, riding in the simulated vehicle. If the agent received external inputs from a subject through the natural interface, then the experimenter would interpret this input to the agent and instruct a set of changes to be made to the driving settings (that the agent will also log as training data). When the agent acted in adaptive mode, changes to settings were also made when the agent proactively predicted human discomfort and the agent decided on the changes to make to the driving parameters based on the learned models.

In the next sections, we describe the three experiments run with the adaptive driving agent. In all, we evaluate the performance of the agent by quantifying the number of human interventions needed to attain a comfortable ride. All data collected included the vehicle dynamics, driving contexts and subjects' inputs if provided.

4 Default Driving Agent

The goal of the first experiment was to evaluate how many manual corrections would be requested by human subjects to adjust a default driving style during a set of simulated rides. Our hypothesis was that humans are different and therefore driving contexts will affect the preferred style of driving. We recruited 30 subjects (17 males and 13 females), ranging in age from 21 and 50 (avg. 33, s.d. 7.02). Participants were told that they were going to take a ride in a simulated taxi in San Francisco for 4 laps driven by a simulated automated driver. All subjects started 2 rides with the default *calm* driving style and 2 more rides with the *aggressive* style as in Table 1 (all runs were balanced). Users interacted with the simulated vehicle through the adaptive intelligent interface (see Fig. 2). When subjects entered any comment (by using the tablet), the experimenter paused the simulation and asked the subject for his actual intention. Then, the experimenter adjusted the driving parameters until a desired changed was achieved (e.g., when the comment was "The slowing down was harsh", the correction included changing the values of the "Brake Smooth" and the "Forward Approach Brake Smooth" values). Any request for driving corrections remained in effect until the end of the current two rounds. A new driving style was implemented in the next two rounds. Fig. 3 shows a total of 633 comments received from the subjects (avg. of 21.1 per participant, s.d., 9.48); with 26.7% of the comments being related to the events we predefined (these comments were provided up to 7 seconds following an event). For example, when a jaywalker crossed the street in front of the car, many participants commented "Less Speed" since the car was perceived to be approaching the walker too fast (despite having enough time to stop).

From the usability questionnaires we collected, we found that 66.66% of the subjects mentioned that the intelligent interface was very easy to mostly easy to use (average score of 3.2 out of 10, the lower the better). The subjects score the match between the driving corrections and their intent high with an average score of 8.1 out of 10 (the higher the better).

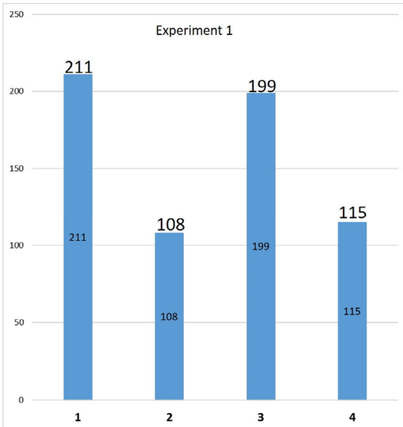

Figure 3: Total Manual Corrections Requested per Rides Laps

Clustering method	Avg. Distance
Clustering by set (1) using K-Means with k=2	6.900
Clustering by set (1) using K-Means with k=3	6.159
Clustering by set (1) and set (2) using K-Means with k=2	8.740
Clustering by set (1) and set (2) using K-Means with k=3	8.416
Clustering by set (1), set (2) and set (3) using K-Means with k=2	9.058
Clustering by set (1), set (2) and set (3) using K-Means with k=3	7.572
Clustering by set (1) and set (3) using K-Means with k=2	7.451
Clustering by set (1) and set (3) using K-Means with k=3	**5.897**
Clustering by set (1) and set (3) using DBSCAN with 2 clusters	7.821
Clustering by set (1) and set (3) using DBSCAN with 3 clusters	7.542

Figure 4: Human Data-Driven Driving Clusters

Parameter	Normal	Calm	Active
Gas Smooth	0.177419	0.548	0.1
Brake Smooth	0.177419	0.64	0.1
Gap Distance	6.048387	8.81	5.625
Gap Time	1.955645	2.365	1.5625
Forward Approach Gas Smooth	0.251613	0.732	0.2
Forward Approach Brake Smooth	0.26129	0.66	0.3
Default Acceleration	23.08468	15.25	25
Maximum Speed	47.58065	40.2	50
Lane Centering	0	0	0
Turn Speed	28.51613	19.68	32

Table 2: Human Data-Driven Driving Style Profiles

5 Human Data-Driven Agent

Our hypothesis was that the number of manual corrections can be reduced when users *choose* what driving style they prefer. Moreover, they chose from styles learned from the **human-data** collected in experiment 1. This data reflected the driving settings that converged towards the end of each round of experiment 1 when the user did not make any further comments. We tested clustering this data to find the densest types of driving settings (i.e., all data points in one cluster that have the shortest distance from the centroid). That is, we scored a clustering method with the standard density measure that computed the average of distances between items inside the cluster from their centroid and then the average among clusters. Let 'm' be the number of clusters, 'x' be a vector pointing to the data point and 'c' be a vector pointing to some centroid. Then the distance 'd_i' to the centroid can be defined as: $\sqrt{(\vec{x} - \vec{c})^2}$ and the density measure is the average of this value over the m clusters. The x vector comprised (1) the driving style parameters at the end of rounds 2 and 4, (2) the number of each type of corrections made, (3) averages of speed, acceleration and jerk. Fig. 4 shows the clustering algorithms tested (K-means and DBSCAN [14,7]) and their corresponding scores (K-means was tested for various values of k: values equal to 2 and 3 are included in the table, larger values were not found to lead to better clustering solutions). Note that the types of corrections provided do not characterize the styles, but the driving dynamics do. DBSCAN was too sensitive to small changes in its parameters; it did not find a balanced division of points to styles. The winning clustering algorithms was K-means (on inputs 1 & 3); this is consistent with results presented in [24]. The novelty is that the styles were learned from human data rides that converged (see Table 2). Normal is situated between the two others. All scores are statistically significant different (tested with ANOVA).

In experiment 2, we recruited 30 new subjects (16 males and 14 females), ranging in age from 23 and 47. The procedure of experiment 2 differs from that in experiment 1 in the users **choosing** the initial (**learned**) driving styles (in experiment 1, the initial driving styles were set as default styles, all participants experienced the same default styles with counterbalanced order). In this experiment, participants were asked to choose a driving style they would prefer in a taxi-like ride. They could choose a calm or sportier style or one in between these to reflect the three styles learned by our clustering algorithm. A reduction of 42% in the number of manual corrections was indeed attained by getting only a total number of 367 corrections from the 30 subjects (i.e., on average, subjects requested 12.2 corrections). We noted that even in a simulated study as we run, different users had different number of comments, meaning some are better with some style chosen and some requested some additional adjustments (some subjects had only a few corrections, while others had as many as 21 comments).

From the usability questionnaires we collected, we found that the subjects gave an average score of 3.2 (out of 10) to the easiness of use of the adaptive interface (the lower the better). The subjects score the match between the driving corrections and their intent high with an average score of 7.9 out of 10 (the higher the better).

6 Adaptive Driving Agent Solution

Our end goal was to show how an adaptive agent can improve human driving-comfort, by adjusting the automated driving settings in real time. We developed such agent that first, learned successfully a model of human discomfort from driving. Then, the agent learned what actual correction should be executed online, once discomfort is predicted. Our hypothesis is that when users choose their preferred initial driving style (from those learned from human data) and further interact with an agent that adapts its driving behavior to their expected preferences, users will be intervening the least, compared to the results in the previous experiments. The same experimental procedure was applied with the adaptive agent version implemented this time.

[Offline] Discomfort Model - We first pre-processed the collected data such that each half a second time frame was associated with the next set of 24 features:

1. Front Distance To - the distance from the car in front
2. #Surrounding Cars - Number of cars in a fixed radius
3. #Surrounding Cars (Adaptive) - Number of cars in a speed-dependent radius
4. Speed - Current speed of the car (Km/h)
5. Acceleration - Current acceleration of the car (km/h2)
6. Avg. Speed - average speed of the car for the last 8 seconds
7. Avg. Acceleration (calculated the same way as above)
8. Lateral Acceleration (m/s2)
9. Longitudinal Acceleration (m/s2)
10. Lateral Jerk (m/s3)
11. Longitudinal Jerk (m/s3)
12. Is Max Speed Reached? (1/0)
13. Driving Behind Another Car? (1/0)
14. – 24. 11 Predefined Events

We chose 11 timeframes (equal to a total of 5.5 seconds) in a moving-window fashion to construct the training instances since users' inputs are not instantaneous. Each instance is classified as *satisfied* (1) or *unsatisfied* (0) to represent the users' comfort from driving with that vector assignments (SMOTE [5] was used to artificially balance the data set). We built classifiers to predict when these features' settings will result in driving-discomfort for the user. Driving discomfort is understood as an event when the participant provides input to make an adjustment to the current driving settings. While the participant does not provide any input to adjust these settings, we assume that the participant is comfortable with the driving style experienced. We measured the quality of these classifiers with the standard Area Under the Curve (AUC) score [12]. Using data from experiment 1 as a training set, we evaluated different prediction models [23]. Random Forest was too simple to capture the dependencies between the driving settings and discomfort. Also, Linear Regression did badly since our data is not linearly separable (due to lack of space the graph is not included). However also the Multi-layer perceptron could not predict well. So, we evaluated networks that can capture the time sequence relationships in the data. CNN resulted as a better predictor than the Long-Short-Term Memory. Table 3 summarizes the AUC scores attained by all predictors tested. Increasing the training set with data from experiment 1 with data collected in experiment 2 (a total of 117,100 data points) and changing the number of filters and their size, improved the accuracy to 95.96% in training and 95.44% in testing with an AUC reaching 0.85. This retrained CNN could anticipate a user's comment by 2.5 seconds.

[Online] Adaptation - Every half second the simulator sent data (time, position, driving settings, acceleration, jerk and predefined events) to the agent. Every 5.5 seconds of simulation, the logs are processed into the 24 features that activate the CNN to predict the current user's driving discomfort (see Fig. 5). If the agent predicts a state of discomfort, then the agent searches its training data set for corrections already executed in situations like the current one (see Fig. 6). The 9 samples closest to the current state are found and the correction with the highest probability (max vote) is chosen. For example: let c_i (i = [1,9]) denote the 9 data points found closest to the current state; then without loss of generality assume: c1-2=More Speed, c3-9=More Sport. Then, with probability of 7/9, *More Sport* will be chosen and *More Speed,* with probability of 2/9. Still, the user can enter their input using the interface at any time (e.g., to reject the correction performed by the agent).

To evaluate the adaptive agent performance, we recruited 34 subjects (19 males and 15 females), age range in 25-52. Fig. 7 shows the results comparing the 3 conditions tested; error bars indicate standard error. We can clearly see that (1) the use of human data-driven driving styles brings about a significant decrease in the number of modifications made by the users, $p < 0.05$ (i.e., 12.2 corrections on average vs. 20.9 when default driving styles were initialized with no user choice). (2) the adaptive agent successfully reduced this number compared to both other conditions in a statistically significant manner, $p < 0.05$ (i.e., only 7.4 corrections on average vs. 12.2 were required for the participants to achieve acceptable driving comfort). The statistical analysis was performed using an ANOVA test followed by post-hoc t-tests comparisons with Bonferroni correction. On average, our adaptive agent performed 19.7 autonomous changes to the driving settings during the ride (s.d. 6.3). On average, users accepted 16.2 of these (accepted means that the agent predicted discomfort and consequently adjusted the driving settings even prior to the participant actually asking for these changes). On average, users corrected the agent only 3.75 times (meaning even when the agent predicted discomfort and adjusted the settings, the user either did not agree with the prediction or with the settings adjustment done automatically). Finally, on average, users gave 3.65 additional corrections (that the agent did not predict discomfort accurately). The number of requests depends on the initial driving style (see Fig. 8) and the time passed from the beginning of the drive.

From the usability questionnaires we collected, we found that the subjects gave an average score of 3.6 (out of 10) to the easiness of use of the adaptive interface (the lower the better). The subjects score the match between the driving corrections and their intent high with an average score of 8.6 out of 10 (the higher the better).

Algorithms	AUC score
Random Forest	0.69
Linear Regression	0.59
Multi-Layer Perceptron	0.69
Long-Short Term Memory	0.67
CNN without acceleration/jerk inputs	0.75
CNN with acceleration/jerk inputs	0.69
CNN with lateral and longitudinal acceleration	0.68
CNN with acceleration/jerk inputs (trained on data from experiments 1 &2)	0.85
CNN trained on Comfort initial driving style	0.68
CNN trained on Normal initial driving style	0.7
CNN trained on Sport initial driving style	0.63

Table 3: Predicting Driving Discomfort

Algorithm 1 ACCA

Require: D: Data set used for training the CNN

1: **while** *simulation active* **do**
2: // Satisfaction prediction
3: **for each** *vector* received from the simulator **do**
4: *vec_norm* ← Normalize(*vector*) // *With the same avg. and s.d. as in D*
5: *prediction* ← CNN_Prediction(*vec_norm*)
6: **if** *prediction* is 'unsatisfied' **then**
7: *correction* ← Get_Correction(*vec_norm*, *D*)
8: **append** *correction* **to** *D*
9: **send** *correction* **to** *simulator*
10: // Manual corrections
11: **for each** *correction* made by the user **do**
12: **append** *correction* **to** *D*

Figure 5: Adaptive Driving Agent Behavior: Discomfort Prediction

Algorithm 2 Get Correction

1: **Inputs:**
vector: Current driving log data
D: User corrections (inc. data from training set)
2: **Initialize:**
h **as** *minimum heap*
correction_candidates **as** *list*
corrections_types **as** *array* **of** *list*
3: **for each** *sample* ∈ *D* **do**
4: *similarity* ← Euclidean_Distance(*sample*, *vector*)
5: **append** <key=*similarity*, value=*sample*> **to** *h*
6: *candidates* ← **pop** 9 nearest neighbors **from** *h*
7: *correction_type* ← **choose** correction type **from** *candidates* such that the probability of each correction type is defined as its ratio among the neighbours
8: *result* ← **apply** *correction_type* **to** *vector* according to the rules
9: **return** *result*

Figure 6: Adaptive Driving Agent Behavior: Discomfort Mitigation

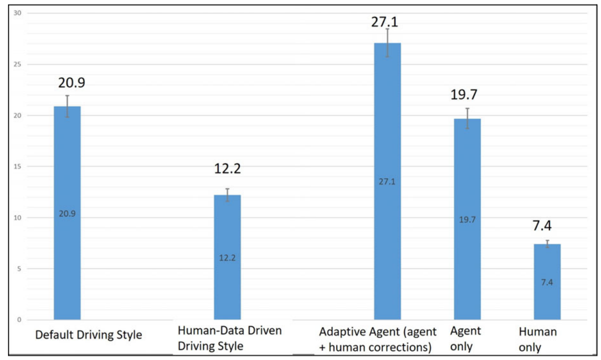

Figure 7: Average number of manual comments in all tests

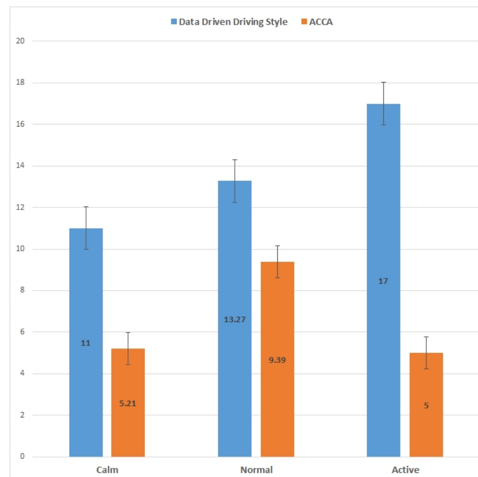

Figure 8: Average number of manual comments per chosen driving style

7 Conclusions

We introduced an automated agent for adapting driving profiles to users and contexts to reduce human driving discomfort. Human discomfort was expressed by the participants each time they provided input to adjust their current driving settings. Our agent decreased the number of manual interventions required to correct driving settings. We also introduced a new approach for human modeling, based on the Convolutional Neural Network (CNN) trained on data obtained through an adaptive intuitive interface. Using K-means clustering and the CNN, we showed that this algorithm can be used in training supervised deep network models. We conclude that we can successfully integrate human models of preferences into the automated control systems to improve their utilization and effectiveness. Instead of a unilateral interaction between a driver and an automated vehicle where the user just operates this machine, we created a bi lateral interaction where the machine is aware of its user. The machine learns from users' inputs, it predicts users' needs and proactively adjusts its own settings to increase user satisfaction and acceptance. The complete AI agent system together with the novel adaptive interface were tested and evaluated successfully through 3 user studies covering a total of 94 human participants in a simulated set up of automated driving scenes.

ACKNOWLEDGMENTS

We thank Shlomi Zigart for the interface design.

REFERENCES

[1] N. E. Assmann, et al. "Should my car drive as i do? what kind of driving style do drivers prefer for the design of automated driving functions?" 17 Braunschweiger Symp. AAET, ITS automotive nord, pages 185–204, 2016.

[2] M. Beggiato et al. "Driving comfort, enjoyment and acceptance of automated driving effects of drivers' age and driving style familiarity". *Ergonomics*, 61(8), 2018.

[3] I. Bae et al. "Self-driving like a human driver instead of a Robocar". *In 3rd International Conf. on Electric Vehicle, Smart Grid and Information Technology*, 2018

[4] Basu et al. "Do You Want Your Autonomous Car to Drive Like You?". *HRI*. Vienna, Austria. 2017.

[5] N.V. Chawla et al. "Smote: synthetic minority over-sampling technique". *Journal of artificial intelligence research*, 16:321–357, 2002.

[6] A. Degani et al. "On sensitivity and holding in automotive systems: The case of the climate control". *Proceedings of the Human Factors & Ergonomics Society Annual Meeting*, volume 60:1906–1910. LA, CA, 2016.

[7] M. Ester et al. "A density-based algorithm for discovering clusters in large spatial databases with noise". *KDD* volume 96, pages 226–231, 1996.

[8] L. Fridman et al. "What can be predicted from six seconds of driver glances?". *Procs. of the CHI Conference on Human Factors in Computing Systems*, pages 2805–2813. ACM, 2017.

[9] X. Geng et al. "A scenario-adaptive driving behavior prediction approach to urban autonomous driving". *Applied Sciences*, 7(4):426, 2017.

[10] J. Guna et al. "Estimation of the driving style based on the users' activity and environment influence". *Sensors*, 1(2404), 2017.

[11] S.J. Hoch and G.F. Loewenstein. "Time-inconsistent preferences and consumer self-control". Journal of consumer research, 17(4),1991.

[12] J. Huang and C. X Ling. "Using auc and accuracy in evaluating learning algorithms". *IEEE Transactions on knowledge and Data Engineering*, 17(3):299–310, 2005.

[13] J. Hwang et al. "Expressive driver-vehicle interface design". *Procs. of Designing Pleasurable Products and Interfaces*, page 19. ACM, 2011.

[14] A. K Jain et al. Algorithms for clustering data, volume 6. Prentice hall Englewood Cliffs, 1988

[15] L. Li et a. "Cognitive cars: A new frontier for adas research". *IEEE Transactions on Intellignt Transportation Systems*, 13(1):395–407, 2012.

[16] G. Mazulla et al. "How drivers' characteristics can affect driving style". *Transportaton Research Procedia*, 27:945–952, 2017.

[17] I. Politis et al. "To beep or not to beep?: Comparing abstract versus language-based ultimodal driver displays". *Procs. of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015.

[18] A. Rosenfeld et al. "Online prediction of exponential decay time series with human-agent application". *Procs. of the 22nd European Conference on Artificial Intelligence*, pages 595–603. 2016.

[19] A. Rosenfeld and S. Kraus. "Predicting human decision-making: From prediction to action". *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(1):1–150, 2018.

[20] A. Rosenfeld et al. "Learning drivers' behavior to improve adaptive cruise control". *Jl. of Intelligent Transportation Systems*, 19(1):18–31, 2015.

[21] K. E Schaefer and E. R Straub. "Will passengers trust driverless vehicles? removing the steering wheel and pedals. In *IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support*, pages 159–165. IEEE, 2016.

[22] C. Senouth et al. "Personalized lane keeping assist strategy: Adaptation to driving style". *IET Control Theory and Applications.* 2018.

[23] S. Shalev-Shwartz and S. Ben-David. "Understanding machine learning: From theory to algorithms". Cambridge university press, 2014.

[24] O. Taubman-Ben-Ari and V. Skvirsky. "The multidimensional driving style inventory a decade later: Review of the literature and re-evaluation of the scale". *Accident Analysis & Prevention*, 93:179–188, 2016.

[25] L. Xu et al. "Establishing style-oriented driver models by imitating human driving behaviors". *IEEE Transactions on Intelligent Transportation Systems*, 16(5):2522–2530, 2015