

MAVERIC: A Data-Driven Approach to Personalized Autonomous Driving

Mariah L. Schrum, Emily Sumner, Matthew C. Gombolay, and Andrew Best

Abstract—Personalization of autonomous vehicles (AVs) may significantly increase acceptance. In particular, we hypothesize that the similarity of an AV’s driving style compared to a user’s driving style, the level of aggressiveness of the driving style, and other subjective factors (e.g., personality) will have a major impact on user’s willingness to use the AV. In this work, we 1) develop a data-driven approach to personalize driving style and calibrate the level of aggressiveness and 2) investigate the subjective factors that impact user preference. Across two human subject studies ($n = 54$), we demonstrate that our approach can mimic the driving styles and tune the level of aggressiveness. Second, we leverage our framework to investigate the factors that impact homophily. We demonstrate that our approach generates driving styles objectively ($p < .001$) and subjectively ($p = .002$) consistent with end-user styles ($p < .001$) and can effectively isolate and modulate a dimension of style (i.e., aggressiveness) ($p < .001$). Furthermore, we find that personality ($p < .001$), perceived similarity ($p < .001$), and high-velocity driving style ($p = .0031$) significantly modulate the effect of homophily.

Index Terms—personalization, autonomous vehicles, human-robot interaction.

I. INTRODUCTION

Driving style is defined as the characteristics of driving related to the judgment and decisions of the driver in a specific situation [15]. Research has shown that driving styles

This paper was produced by the IEEE Publication Technology Group. They are in Piscataway, NJ.

Mariah L. Schrum and Matthew C. Gombolay are with the Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta 30332, USA, (e-mail: mschrum3@gatech.edu; matthew.gombolay@cc.gatech.edu)

Emily Sumner and Andrew Best are with Toyota Research Institute, 4440 El Camino Real, Los Altos, CA 94022, USA, (e-mail: emily.sumner@tri.global; andrew.best@tri.global)

differ greatly amongst individuals [39]. For example, the way in which a driver interacts with other drivers, the level of aggressiveness that a driver exhibits, and tendency to commit traffic violations are characteristics that define an individual’s unique driving style and these characteristics vary across end-users. Because of these individual differences, when riding in an autonomous vehicle (AV), prior work suggests that end-users’ expectations and preferences for the behavior of the AV will likely be influenced by their own driving style [20], [21], [35]. One-size-fits-all models employed by AVs which ignore driver differences may lead to decreased acceptance [20]. Instead, the driving style of AVs should be personalized to fit the preferences and expectations of individual end-users.

Much of the prior work in optimizing AV driving styles has assumed that, to increase end-user acceptance and trust, AVs should mimic end-users’ unique driving styles [16], [35]. However, even if we are able to personalize an AV’s behavior, not all end-users will necessarily want the AV to drive *exactly* as the end-user drives [6], [43]. In fact, prior work has suggested that some end-users may want an AV to drive more cautiously than they drive [6], [16], [43]. Additionally, factors such as trust and familiarity with AVs and various personality traits may affect preference for driving styles similar to one’s own [11], [14], [31].

Based upon evidence from prior work [11], [14], [16], [31], we hypothesize that the optimal driving style for an AV is a function of both the end-user’s own driving style and various subjective factors, such as personality. Therefore, to optimize driving style, an AV must be capable of learning about an end-user’s own driving style as well as modulating this style based upon relevant end-user characteristics. We identify two key shortcomings in prior work. First, prior approaches that



Fig. 1: This figure shows the 6-DOF driving simulator that we employ in our study. Using this simulator, we demonstrate that MAVERIC is capable of mimicking driving style as well as modulating the level of aggression. Additionally, we show that personality, perceived similarity, and high-velocity driving style impact the effect of homophily.

are capable of modulating the driving style of an AV do not take into account the end-user's own driving style [3], [17], despite prior work indicating that end-user's own driving style is an important predictor of the optimal driving style. Second, prior works which mimic an end-user's driving style [8], [25] are not capable of modulating various aspects of driving style to account for end-users who do not wish the driving style to exactly mirror their own.

Our goal in this work is to overcome these limitations and develop a framework that can mimic an end-user's own driving style while also having the flexibility to modulate the level of aggressiveness to match the preference of the end-user. Mimicking the end-user's own driving style provides us with a good initial guess of the optimal driving style of the end-user. However, because not all end-users may want an AV to exactly mimic their style, our framework can also modulate or fine-tune the aggressiveness of the driving. We choose to modulate aggressiveness because prior work has indicated that this dimension has a large impact on end-user preference [6], [16], [43]. In our work, we define optimal driving style as the style which maximizes preference of the end-user.

We additionally leverage our framework to investigate the effect of homophily (i.e., preference for a style similar to one's own) and identify the subjective characteristics (i.e., personality traits, trust, etc.) that are predictive of preference. We do this by modulating the driving style to be more or less similar to that of the end-user. Our framework enables us to both study preference in the context of homophily as well as mimic and modulate style to ultimately match end-user preference.

We introduce Manipulating Autonomous Vehicle Embedding Region for Individuals' Comfort (MAVERIC). By observing the driving of an end-user, MAVERIC learns a high-level model via a neural network architecture that predicts personalized control parameters for low-level controllers to mimic the driving style of the end-user. Simultaneously, MAVERIC learns a personalized embedding representing the driving style of an end-user. By shifting the personalized embedding along the gradient of aggressiveness, MAVERIC is capable of tuning the AV driving style to be more aggressive or cautious while maintaining other characteristics such as minimum headway distance.

In a human-subjects study, we demonstrate the ability of MAVERIC to effectively mimic an individual's driving style as well as modulate aggressiveness *with respect to one's own driving style* while maintaining other aspects of driving style. MAVERIC's ability to modulate aggressiveness enables us to additionally study the effect of homophily and determine which subjective factors impact an end-user's preference for a style that differs in aggressiveness compared to their own. We demonstrate that preferred driving style is related both to one's own style as well as personality traits, perceived similarity, and self-reported high-velocity driving. We validate our approach in a highway domain and in future work we aim to investigate our approach in more complex and varied domains.

In this work we contribute the following:

- 1) We formulate MAVERIC, a novel framework to personalize driving style and modulate aggressiveness while




Fig. 2: This figure shows our domain of light traffic and associated state information. $v_t^{(ev)}$ is the velocity of the ego at time t , $v_t^{(lv)}$ the velocity of the leading vehicle, $d_t^{(x)}$ the distance between the leading vehicle and ego in the x direction at time t , and $d_t^{(y)}$ the distance in y at time t .

maintaining other aspects of driving style.

- 2) We demonstrate that MAVERIC can closely match an end-user's driving style ($p < .001$) as well as produce more aggressive ($p < .001$) and more cautious ($p < .001$) driving in a high-fidelity driving simulator.
- 3) We find that personality ($p < .001$), perceived similarity ($p < .001$), and high-velocity driving style ($p = .0031$) significantly impact the effect of homophily.

II. RELATED WORK

Driving style is defined as the characteristics of driving related to the judgment and decisions of the driver in a specific situation [15]. Prior work has shown that human drivers exhibit a vast array of different driving styles, and various approaches have been proposed to categorize driving style. For example, Taubman-Ben-Ari et al. divide driving style into four different categories: risky, anxious, dissociative, and distress reduction driving [37]. Other work categorizes types of driving styles into aggressive or defensive [28]. Given the diverse range of driving styles, previous studies indicate that end-users expect their AVs to drive in a specific manner that is likely related to their own driving style [6]. To meet the expectations of human end-users, AVs must be capable of learning about their end-users and personalizing their driving styles accordingly.

Because prior work has suggested that driving style has a large impact on preference and that preference varies across end-user, we develop a framework that is capable of adjusting driving style to meet end-user preferences.

A. Aggressive Driving Style

A common way to categorize driving style is via the level of aggression. Aggression can be measured objectively or subjectively [7], [37]. For example, Bellem et al. quantify driving style via objective metrics including jerk and headway distance [7]. To gain an understanding of a driver's view of their own aggressive style, Harris and Norman developed the Aggressive Driving Behavior Scale [19].

Prior work has investigated the impact of the level of aggression of an AV's driving style on end-user acceptance. Ekman et al. conducted a Wizard-of-Oz study using a Volvo vehicle with a professional driver demonstrating the various driving styles. The authors compared defensive and aggressive driving styles found that a defensive driving style produced higher trust scores [16]. Similar work by Basu et al., Yusof et al.,

and Karlsson et al. found evidence supporting the important impact of the level of aggression on end-user preferences [6], [24], [43]. Because of the large body of literature showing the importance of aggressive driving style, aggressiveness of an AV should be taken into account when designing the optimal AV driving style.

Due to the large impact that aggressiveness has on preference, our approach focuses on modulating the level of aggressiveness to meet end-user preferences.

B. Mimicking Driving Style for Improved End-User Experience

To increase trust, likeability and overall acceptance of the AV there is a need for personalized control frameworks which are capable of adapting to the driving style of individual end-users. Prior work has introduced a variety of approaches to adapt AVs to meet end-user needs and expectations. Many of these approaches aim to mimic an end-user's own driving style. For example, Kuderer et al. utilized an inverse reinforcement learning (IRL) approach to produce personalized AV behavior via a learned cost function [25]. The authors employ a feature-based reward function learned from human data to mimic the driving style of the end-user. They evaluate their approach by comparing average acceleration and jerk over trajectories. Based upon this evaluation, the authors show that their approach captures distinct driving styles. While this method was capable of learning distinct driving styles for different users, the authors did not evaluate their approach in a human subjects study and their approach is not capable of modulating style.

Other work investigates personalization of specific aspects of driving to mimic that of an end-user [8], [18]. For example, Bolduc et al. developed an approach to match driver's style for adaptive cruise control [8]. The authors extract parameters from the end-user's own driving and utilize these parameters to inform the cruise control. Feng and Yan explored personalization of lane changes via a support vector machine [18]. By collecting data from lane changes of end-users, the authors train a personalized support vector machine (SVM) to mimic the lane changing style of the end-user. While these approaches have demonstrated that personalizing an end-user's own driving style can lead to increase acceptance and trust in the AV, the question still remains as to whether or not mimicry is the optimal strategy for an AV.

Prior works that learn to mimic the driving style of the end-user via IRL or supervised approaches presuppose that mimicry is the only relevant objective for personalization. In our work, we aim to learn a more expressive representation of end-user driving style that allows us to go beyond simple mimicry.

C. Should We Mimic End-User Driving Styles?

Despite the many approaches that have been developed for mimicking the driving styles of an end-user, prior work suggests that end-users may not want an AV to drive *exactly* as they drive. Instead various latent factors may influence a driver's preference for a specific driving style that may differ

from their own driving style [6], [16], [43]. For example, A Wizard-of-Oz study conducted by Yusof et al. found that many end-users prefer a more defensive AV driving style compared to their own. In this work, the authors conducted a study using an Audi test vehicle to realistically simulate an AV. Based upon a self-reported questionnaire, participants were categorized as either defensive or assertive drivers. Then each participant experienced a defensive and assertive AV driving style. The authors found that preference for an assertive or defensive driving style depended on the driver's own style and that aggressive drivers prefer more defensive AVs [43].

Basu et al. conducted a study investigating preference for driving styles both similar and different from one's own [6]. The authors investigated a condition intended to mimic the participant's driving style, an aggressive driving condition, a defensive driving condition, and a distractor condition. Basu et al. evaluated their approach via summary statistics that describe the minimum headway distance, mean headway time, distance headway merge back, and velocity. They compare these metrics across conditions to answer their research questions. We utilize a similar set of metrics to evaluate our approach.

More than half of the participants preferred a driving style different from their own. The largest predictor for preference was *perceived* similarity suggesting that participants did not want the AV to drive as they drive, but instead preferred the AV to drive like the end-user *thinks* they drive. This finding suggests that, because humans often lack introspection and have a poor perception of their own driving style, their preferred style often differs from their actual driving style [6]. Based on these results, the authors suggest that we can not simply rely on mimicry and instead, must also account for other end-user characteristics to determine the optimal driving style.

In keeping with prior work [6], we posit that viewing mimicry as the sole objective when optimizing driving style limits the potential for personalization. Our work aims to overcome this limitation by enabling driving style to be modulated based upon latent factors.

D. Modulating Driving Style for Improved End-User Experience

Rather than mimic the driving style of the end-user, Ling et al. introduced a method to adapt driving style online based on the emotional responses of passengers [27]. This approach utilizes EEG signals to analyze the emotions of the passenger and then employs this information to automatically adapt the driving style of the vehicle to match the emotional state.

Prior work has introduced several classic control approaches which allow for tuning of the level of AV aggressiveness. For example, Eriksson and Svensson [17] introduced a linear quadratic controller for tuning the driving of an AV to optimize ride quality. Bae et al. [3] introduced an approach which allows the user to specify the desired parameters for a controller to adjust the driving style of the vehicle with respect to acceleration and jerk. Bae et al. evaluate their approach via a

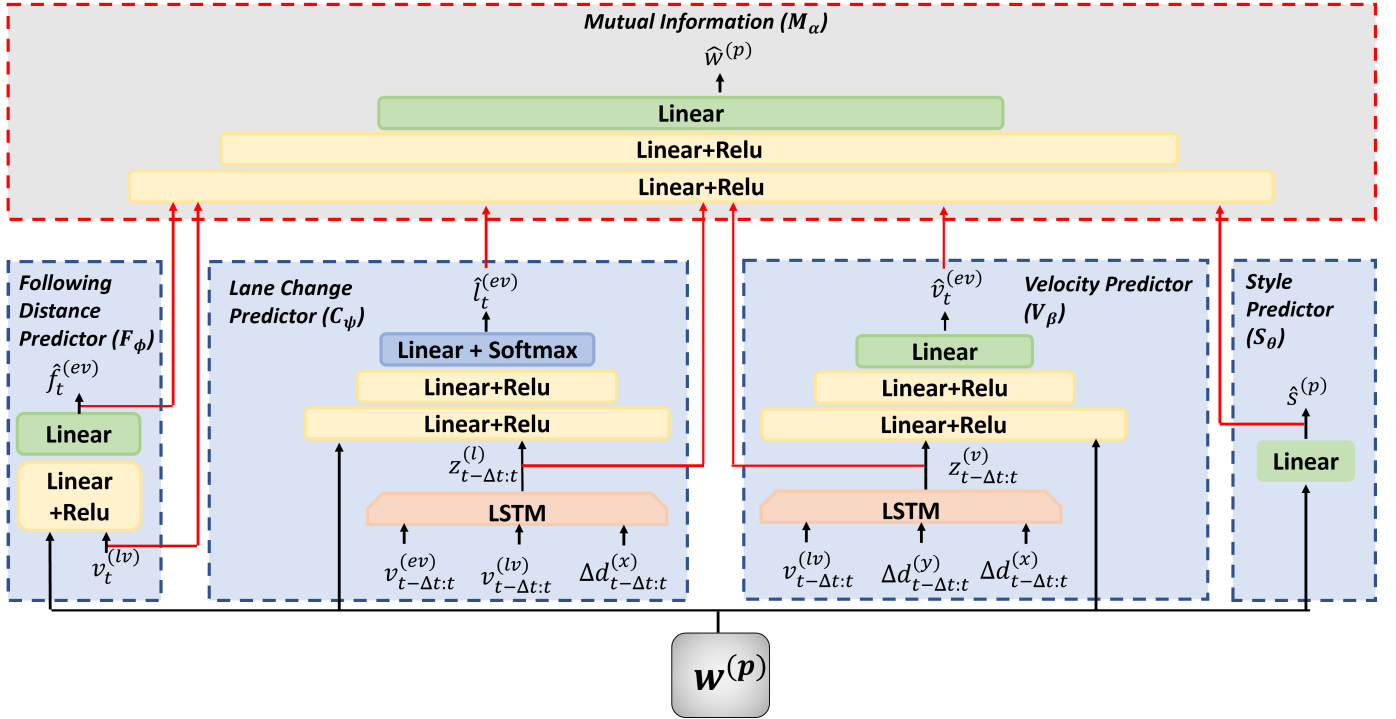


Fig. 3: This figure shows our network architecture. F_ϕ predicts the following distance. C_ψ predicts when a lane change should occur for the ego vehicle. V_β outputs the velocity of the ego vehicle. S_θ is the style predictor subnetwork which predicts the subjective aggressive style of the participant from the personalized embedding, $w^{(p)}$. $v_{t-\Delta t:t}^{(ev)}$ is the ego velocity and $v_{t-\Delta t:t}^{(lv)}$ is the velocity of the lead vehicle from time $t - \Delta t$ to t . $d_{t-\Delta t:t}^{(x)}$ is the distance between the ego and leading vehicle in x and $d_{t-\Delta t:t}^{(y)}$ is the distance in y . $\hat{w}^{(p)}$ is the estimate of the participant’s personalized embedding sampled from the approximate posterior defined by M_α .

novel driving preference metric to capture end-user preference for the various driving styles. However, these approaches do not consider tuning driving style with respect to the end-user’s own style and may require expert knowledge of vehicle dynamics and complex control parameters to determine the correct parameter settings.

While prior work has investigated mimicking driving style and tuning style via controllers, no prior work has created an architecture that can modulate driving style with respect to an end-user’s own style. Yet, prior work provides evidence that this functionality is important for optimizing driving style for an end-user [6], [43]. Additionally, prior work has not extensively investigated the relationship between subjective factors and preference for styles similar to one’s own. In our work, we seek to fill these gaps by proposing an approach capable of producing more or less aggressive behavior with respect to the end-user’s own driving style. Lastly, we conduct a thorough investigation into the factors that impact the effect of homophily.

III. METHODOLOGY

In the following section we provide an overview of MAVERIC as shown in Fig. 4. We discuss our architecture (Fig. 3) and how we endow our framework with the ability to both mimic an end-user’s driving style as well as modulate aggressiveness. Fig. 2 depicts the state information relevant to our architecture.

By maximizing mutual information between a learned personalized embedding and the end-user’s driving behavior, our

framework captures information about the end-user’s driving style and utilizes this information to mimic their driving style (Section III-B). We train MAVERIC on a set of training participants (Section III-B) to learn the network parameters that map the learned personalized embedding and relevant state information (e.g., distance from leading vehicle) to high level control parameters (e.g., velocity). To mimic the driving style of a new driver, we freeze the network parameters and simply learn the new driver’s personalized embedding based upon data collected from observing the driver (Section III-B). To tune the level of aggressiveness, we determine the gradient of aggression within embedding space and shift the personalized embedding along this gradient (Section III-D). A detailed description of each of these capabilities is provided in the following sections.

A. Network Architecture

Our network architecture is depicted in Fig. 3. Our network simultaneously learns the high-level parameters (e.g., timing of lane change) of low level controllers (e.g., lane change controller) and the personalized embedding, $w^{(p)}$, representing the driving style of an individual, p . See Section III-E for details on the low-level controllers. Our network is composed of five subnetworks: the Following Distance Predictor, F_ϕ , Lane Change Predictor, C_ψ , Velocity Predictor, V_β , Style Predictor, S_θ , and Mutual Information, M_α with parameters ϕ , ψ , β , θ , and α . Because humans rely on historical information to make decisions, to better model human decision making, we

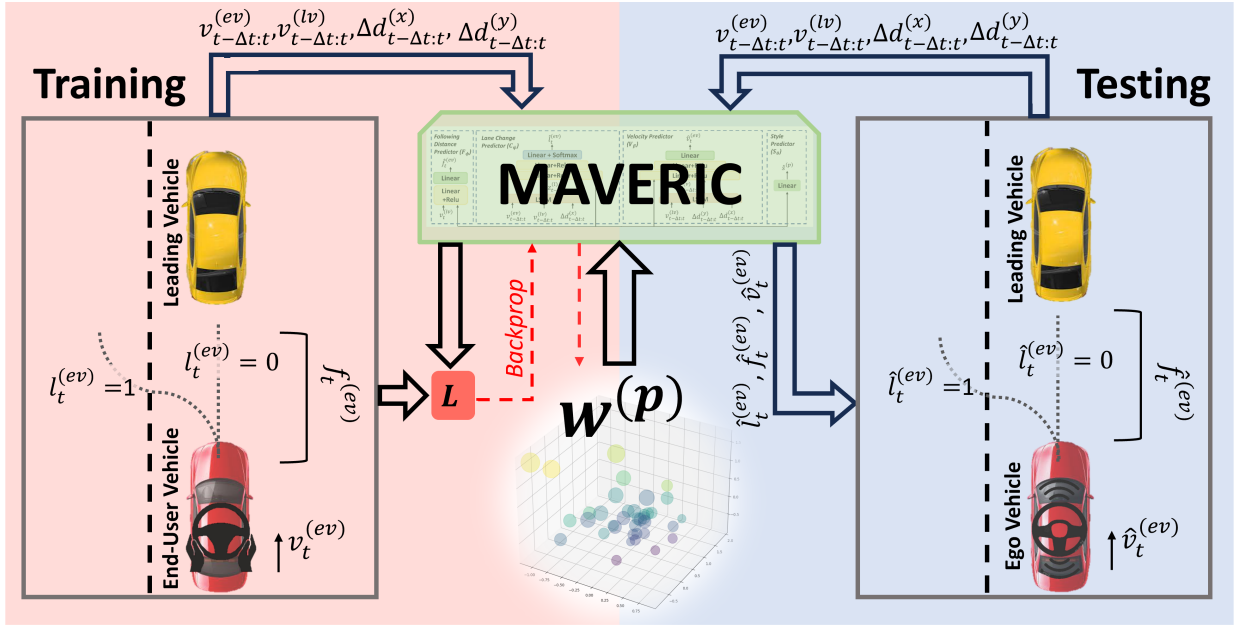


Fig. 4: This figure shows an overview of our MAVERIC training and testing framework. During training, we collect information about the end-user’s driving data including their velocity ($v_t^{(ev)}$) following distance ($f_t^{(ev)}$), and lane changes ($l_t^{(ev)}$). We use this data to learn the personalized embedding, $w^{(p)}$, via loss, L , describing the driving style of the end-user. During testing, this personalized embedding along with relevant state information are fed into the MAVERIC framework to predict the high-level control parameters ($\hat{v}_t^{(ev)}$, $\hat{f}_t^{(ev)}$, and $\hat{l}_t^{(ev)}$) to mimic the end-user’s driving style. A more detailed depiction of the MAVERIC framework is shown in Fig. 3.

employ an LSTM based architectures in our subnetworks when interaction history may be informative for predictions [40]. We utilize LSTMs instead of transformers because transformers would likely require prohibitive amounts of data to achieve similar performance. Each of the subnetworks relies on the personalized embedding, $w^{(p)}$ to inform the predictions.

Following Distance Predictor (F_ϕ)

- Inputs: personalized embedding, $w^{(p)}$, and the velocity of the lead vehicle, $v_t^{(lv)}$.
- Outputs: desired following distance, $\hat{f}_t^{(ev)}$, between the ego vehicle and the lead vehicle.
- Layers: fully-connected with ReLU activations.
- Loss Function: mean-squared error loss defined as $L_1(\phi, w^{(p)}) = \frac{1}{N} \sum_t \|f_t^{(ev)} - \hat{f}_t^{(ev)}\|_2^2$, where $f_t^{(ev)}$ is the ground truth following distance of the end-user.

Lane Change Predictor (C_ψ)

- Inputs: personalized embedding, $w^{(p)}$, ego velocity, $v_t^{(ev)}$, the velocity of the lead vehicle, $v_t^{(lv)}$, and x-distance between the ego and lead vehicle, $d_{t-\Delta t:t}^{(x)}$, from time $t - \Delta t$ to t .
- Outputs: $\hat{l}_t^{(ev)}$, i.e. when a lane change should occur.
- Layers: fully-connected with ReLU and softmax activation.
- Loss Function: cross-entropy loss defined as $L_2(\psi, w^{(p)}) = -\frac{1}{N} \sum_t l_t^{(ev)} \log \hat{l}_t^{(ev)}$ where $l_t^{(ev)}$ is a binary variable indicating a lane change.

Velocity Predictor (V_β)

- Inputs: personalized embedding, $w^{(p)}$, velocity of the lead vehicle, $v_t^{(lv)}$, y-distance, $d_{t-\Delta t:t}^{(y)}$, and x-distance,

$d_{t-\Delta t:t}^{(x)}$, between the ego and lead vehicle, from time $t - \Delta t$ to t .

- Outputs: predicted speed of the ego, $\hat{v}_t^{(ev)}$, at time t .
- Layers: fully-connected with ReLU activations.
- Loss Function: mean-squared error loss defined as $L_3(\beta, w^{(p)}) = \frac{1}{N} \sum_t \|v_t^{(ev)} - \hat{v}_t^{(ev)}\|_2^2$ where $v_t^{(ev)}$ is the velocity of the end-user.

Style Predictor (S_θ)

- Inputs: personalized embedding, $w^{(p)}$.
- Outputs: subjective aggressiveness, $\hat{s}^{(p)}$, of participant, p .
- Layers: fully-connected.
- Loss Function: $L_4(\theta, w^{(p)}) = \frac{1}{N} \sum_p \|s^{(p)} - \hat{s}^{(p)}\|_2^2$ where $s^{(p)}$ is the subjective aggressive style of the end-user as self-reported via a questionnaire. In Section III-D, we discuss the importance of this subnetwork and how we obtain $s^{(p)}$.

Mutual Information (M_α):

Our goal is to learn a representation of driving style that, when observed, most reduces uncertainty about the end-user’s driving style. By doing so, we ensure that our embedding learns salient features of driving style. This objective can be achieved by maximizing mutual information between the representation (i.e., personalized embedding) and the data describing the driving style. This mutual information term is difficult to optimize in practice because it requires access to an inaccessible posterior for the probability of the embedding given the observation history and driving actions. To address this problem, Barber and Agakov [4] lower bound the mutual information via variational information maximization, which utilizes a synthetic distribution to approximate the posterior.

This posterior is trained by minimizing the MSE between $w^{(p)}$ and $\hat{w}^{(p)}$. Following prior work, we employ variational information maximization and define an auxiliary distribution, M_α , to approximate this posterior [30], [32].

- Inputs: $v_t^{(lv)}$, encodings of the relevant time series state information $z_{t-\Delta t:t}^{(l)}$ and $z_{t-\Delta t:t}^{(v)}$, and the outputs of each of the other subnetworks, $\hat{l}_t^{(ev)}$, $\hat{v}_t^{(ev)}$, $\hat{f}_t^{(ev)}$, and $\hat{s}^{(p)}$.
- Outputs: $\hat{w}^{(p)} \sim \mathcal{N}(\mu^{(p)}, \sigma^{(p)})$.
- Layers: fully-connected layers with ReLU activations.
- Loss Function: mean-squared error loss defined as $L_5(\alpha, w^{(p)}) = \frac{1}{N} \sum_p \|w^{(p)} - \hat{w}^{(p)}\|_2^2$.

With this setup, we train $w^{(p)}$ to capture salient information about an end-user’s driving style by maximizing a lower bound, $L_I(F_\phi, C_\psi, V_\beta, S_\theta, M_\alpha)$, on mutual information. Eq. 1 shows the lower bound on mutual information as derived in Barber and Agakov et al. [4]. \mathcal{X} represents the vector containing the relevant state parameter ($v_t^{(lv)}$, $z_{t-\Delta t:t}^{(l)}$, and $z_{t-\Delta t:t}^{(v)}$). \mathcal{P} represents the vector containing the outputs of the subnetworks ($\hat{f}_t^{(ev)}$, $\hat{l}_t^{(ev)}$, $\hat{v}_t^{(ev)}$, and $\hat{s}^{(p)}$). \mathcal{X} and \mathcal{P} are used to recover the distribution, $\hat{w}^{(p)} \sim \mathcal{N}(\mu^{(p)}, \sigma^{(p)})$. $\hat{w}^{(p)}$ is the estimate of the participant’s personalized embedding sampled from the approximate posterior defined by M_α .

$$I(w^{(p)}; \mathcal{X}, \mathcal{P}) = H(w^{(p)}) - H(w^{(p)} | \mathcal{X}, \mathcal{P}) \geq \mathbb{E}[\log(M_\alpha(w^{(p)} | \mathcal{X}, \mathcal{P}))] + H(w^{(p)}) = L_I(F_\phi, C_\psi, V_\beta, S_\theta, M_\alpha) \quad (1)$$

B. Training Procedure

Alg. 1 shows the training procedure. We train our MAVERIC architecture to minimize the five loss functions, $L_1(\phi, w^{(p)})$, $L_2(\psi, w^{(p)})$, $L_3(\beta, w^{(p)})$, $L_4(\theta, w^{(p)})$, and $L_5(\alpha, w^{(p)})$. Loss functions L_1 through L_4 are utilized to train the four predictor subnetworks. $L_5(\alpha)$ minimizes the MSE between the embedding $\hat{w}^{(p)}$ sampled from the approximate posterior and the true embedding, $w^{(p)}$ ($L_5(\alpha, w^{(p)}) = \frac{1}{N} \sum_p \|w^{(p)} - \hat{w}^{(p)}\|_2^2$). This is equivalent to maximizing the log likelihood of the posterior represented by M_α [30].

We initially train our architecture on a set of data gathered from a large distribution of drivers (which we refer to as training the participants). We note that we initialize $w^{(p)}$ (Alg. 1, Line 9) for each of these drivers by sampling from the prior, $w^{(p)} \sim U(0, 1)$ [30]. The sum of these five losses (Eq. 2) is then backpropagated through each of the subnetworks and $w^{(p)}$ to simultaneously learn the network parameters $\phi, \psi, \beta, \theta$, and α and the personalized embedding representing each driving style (Alg. 1, Line 5). To apply the framework to a new end-user, we first freeze the network parameters (Alg. 1, Line 6). Then we collect data of this new end-user’s driving style and learn the personalized embedding (initialized to the mean of the training participants’ embeddings) describing the new end-user’s driving style (Alg. 1, Lines 9-11). Intuitively, this means that, for a new end-user, we are learning where this end-user’s driving style falls within the distribution of training participants.

Algorithm 1 MAVERIC Procedure (Details about the human subject study can be found in Section IV-C. Interactions with subjects are italicized.)

- 1: **for** q in training participants **do**
- 2: *Collect pre-study survey data from q*
- 3: *Collect driving data from q*
- 4: **end for**
- 5: Perform gradient descent on $\phi, \theta, \psi, \beta, \alpha$, and w until convergence (Eq. 2)
- 6: Freeze $\phi, \theta, \psi, \beta$, and α
- 7: **for** p in test participants **do**
- 8: *Collect pre-study survey data from p*
- 9: Initialize $w^{(p)} \leftarrow \frac{1}{N} \sum_{i=0}^N w^{(i)}$ where $w^{(i)}$ is the embedding of training participant, i
- 10: *Collect driving data from p*
- 11: Perform gradient descent on $w^{(p)}$ until convergence
- 12: *Present study conditions {Mimic, Aggressive, Cautious, and Perpendicular} in randomized order.*
- 13: **for** c in conditions **do**
- 14: Shift embedding as described in Section IV-D
- 15: *Present driving style, c , to participant, p*
- 16: *Collect post-trial survey data*
- 17: **end for**
- 18: **end for**

$$L = \frac{1}{N} \sum_t \left(\|f_t^{(ev)} - \hat{f}_t^{(ev)}\|_2^2 - l_t^{(ev)} \log \hat{l}_t^{(ev)} + \|v_t^{(ev)} - \hat{v}_t^{(ev)}\|_2^2 \right) + \frac{1}{N} \sum_p \left(\|s^{(p)} - \hat{s}^{(p)}\|_2^2 + \|w^{(p)} - \hat{w}^{(p)}\|_2^2 \right) \quad (2)$$

C. Training Data

To learn the network parameters, $\phi, \psi, \beta, \theta$, and α , we collect 10 minutes of driving data at 4hz from 38 examples of driving data (91,200 total samples). We collect data from a diverse set of participants (See Section IV for more details) to cover the distribution of potential driving styles and thereby decrease issues related to covariate shift when MAVERIC is applied to a new participant. To learn the personalized embedding, $w^{(p)}$, of a new driver, we collect 10 minutes of driving data at 4 hz (2400 samples per participant). We find we can represent the driving style of an end-user via a 3-dimensional vector and that increasing the size of the embedding did not significantly improve the accuracy of predictions. Because, for a new-end user, MAVERIC only needs to learn a 3-dimensional personalized embedding rather than learn a large number of neural network parameters, we are able to accurately model the driving style of a driver with relatively few samples. We found that collecting 10

minutes of training data per end-user (2,400 samples) was sufficient to learn driving style within the context of our study setup. However, we hypothesize that more training data will be required to sufficiently learn driving style in more complex scenarios that are common in the real world.

D. Modulating Aggressiveness

We designed MAVERIC to be capable of both matching driving styles of individuals and modulating aggressiveness with respect to an individual’s driving style. Because MAVERIC learns a latent embedding space, we aim to create a dimension of aggressiveness within the embedding space, allowing us to shift an embedding along that dimension and modulate aggressiveness, while keeping other driving characteristics unrelated to aggressiveness constant. To achieve this objective, we add an additional signal when learning the embedding space. Specifically, we add a network head, S_θ , composed of a fully-connected layer which takes as input the personalized embedding, $w^{(p)}$. S_θ is trained to predict the subjective aggressive driving style of the user as measured by the participant’s response to the Aggressive Driving Behavior (ADB) scale [19]. By doing so, we create an aggressive dimension within the embedding space. We then move along the gradient of aggressiveness ($\nabla_w S_\theta$) to produce a more or less aggressive driving style as shown in Fig. 5.

While driving style has multiple dimensions [37], we focus on the aggressive dimension, as prior work has shown that this dimension has a large impact on end-user preference [6], [16], [43]. Other characteristics of driving could be modulated by following a similar procedure. While we acknowledge that the ADB scale is a noisy metric, as discussed in Section V, our results demonstrate that our method can effectively produce more and less aggressive behavior.

E. Low Level Controllers

MAVERIC learns the parameters for low-level controllers (e.g., velocity, timing of lane change, etc.) rather than directly learning the low-level controls (i.e., throttle and steering) to enable safety constraints and to account for unexpected or dangerous behavior that could be produced by the network. For example, by learning the desired following distance for an end-user and utilizing an adaptive cruise controller to maintain this distance, we can ensure that the following distance remains safe. Additionally, by predicting when a lane change should occur via the neural network and utilizing a low-level controller to execute the lane change, we ensure consistent and smooth lane changes. Furthermore, this hierarchical method of learning and control has been shown to produce better results in prior work [29].

We utilized the specific controllers described below because they have been shown to be robust and produce desired behavior in prior work [34]. However, our approach is agnostic to the type of low level controller and these controllers could be exchanged for other controllers.

Lane Change Controller: Our lane change controller is based on a Stanley controller [34] and follows a Bezier curve [2]. We compute the Bezier curve based upon the desired

distance (selected to produce natural behavior) to complete the lane change, while ensuring that the ego vehicle will not collide with the leading vehicle. The lane change controller executes a lane change when $\hat{l}_t^{(ev)} > \delta$.

Velocity Controller: We utilize a proportional and integral (PI) controller to maintain the desired velocity, $\hat{v}_t^{(ev)}$, of the ego vehicle as predicted by the neural network.

Following Distance Controller: When the distance between the ego and leading vehicle falls below threshold, λ , we switch from the velocity controller to the following distance controller. The following distance controller is a PI controller that minimizes both the error between the desired following distance, $\hat{f}_t^{(ev)}$, as predicted by the neural network and the difference in speed between the ego and leading vehicle subject to safety constraints on following distance.

IV. HUMAN SUBJECTS STUDIES

We conducted two human subjects studies: A Model Training Study (Study 1) and a Model Testing Study (Study 2). In Study 1, we collect data (91200 data samples in total) from 30 participants to train MAVERIC and learn θ , ϕ , ψ , α , and β , and the participants’ embeddings. We then freeze these parameters, and in Study 2, we collect driving data from 24 participants to learn their embeddings. Then, to study the effect of homophily, each participant experiences the four AV conditions described in Section IV-D. The study design is fully counterbalanced. Research was approved by an Institutional Review Board (protocol #20221727).

A. Driving Simulator

To evaluate MAVERIC, we utilize a high-fidelity, research grade driving simulator. The simulator (Fig. 1) is an immersive installation featuring a 6-DOF platform capable of emulating the motion of a vehicle. The simulator software is based on CARLA [13], ROS2, and Unreal Engine. The simulator provides high-fidelity state-of-the-art graphics and 120 hz refresh rate. The cabin is equipped with a standard digital speedometer, side and rear-view mirrors, a consumer-grade racing wheel and pedal assembly, and digital shifter. The center-console is a tablet which is used to collect subject responses after each trial. Additional details on the simulator can be found at Medium¹. The scenario environment is a sparse two-lane highway of 50km and standard lane-width. Road-side decorations are omitted to reduce distraction and mitigate motion-sickness. Additional details on the scenario setup can be found in Section IV-C.

B. Participants

Model Training Study (Study 1): We recruited 30 participants (Mean age 35.4; 27% Female) via word of mouth and mailing lists. Four of the participants were professional drivers who demonstrated aggressive, cautious, and their own driving style to ensure that we train on a wide distribution of styles. In total, we collected 38 data points representing various driving styles.

¹<https://medium.com/toyota-research/847f36ea103e>

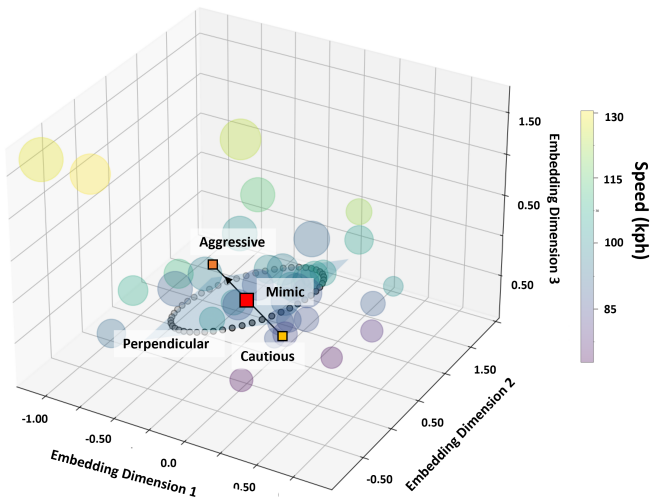


Fig. 5: This figure shows the learned embedding space. The size of the points represents the subjective aggressive style of the participant and color represents the average velocity. The black line shows the vector of the aggressive gradient. The red square represents a candidate learned embedding of a test participant. We shift the embedding along the gradient to increase (orange square) or decrease (yellow square) the ADB [19] score by 15 points to produce behavior for the aggressive and cautious conditions respectively. We randomly sample from the gray points to produce the Perpendicular behavior.

Model Testing Study (Study 2): Study 2 was run with two different populations of participants to increase diversity. For the *internal* study, 12 subjects were recruited internally (Mean age 34.42; 33.4% Female). For the *external* study, 12 subjects (Mean age 43.92; 41.7% Female) were recruited from the general public via Fieldwork recruiting. External participants were compensated \$250. The populations are analyzed as a collapsed dataset because the procedure was identical.

C. Procedure

We investigate personalization of driving styles in the domain of light traffic on a two-lane divided highway with both lanes going in the same direction (Fig. 2). We chose to test our framework in a domain in which participants would be forced to make various decisions and judgements and expose their driving style. At the same time, our goal is for the domain to not be so dynamic and complex that it would be difficult to compare driving styles across individuals and properly evaluate our framework’s ability to capture diverse driving styles. Alg. 1 shows the study procedures.

Study 1: Participants first complete pre-study surveys (Alg. 1, Line 2) to collect information about demographics and attitude towards AVs (Section IV-E). Participants complete a practice session to familiarize themselves with the vehicle controls and domain. Participants control the vehicle and demonstrate their driving style for 10 minutes (Alg. 1, Line 3). Their task is to drive as they would in their own vehicle. They are instructed to maintain the speed they would typically drive if the speed limit is 55mph and to pass other vehicles when they feel it is appropriate. In this domain, participants encounter vehicles in the same lane (lead vehicles) and in the adjacent lane (off-lane vehicles). Participants must make decisions about changing lanes, following distance, and

velocity. The speed of the lead vehicles is randomly selected without replacement from the set $\{0.85v_e, 0.9v_e, 0.97v_e, 0.9s, s, 1.1s\}$ where v_e is the ego target speed and s the posted speed (55mph). These speeds ensure consistency across participants but also ensures that a portion of the leading vehicles are slower than the ego, thus forcing the participant to make a decision about changing lanes. We then use the collected driving data to learn the network parameters and the training participants’ personalized embeddings (Alg. 1, Line 5).

Study 2: In the testing study, we freeze the network parameters, ϕ , θ , ψ , β , and α learned from Study 1 data (Alg. 1, Line 6). Participants fill out the pre-study surveys (Alg. 1, Line 8), complete a practice round, and then drive the vehicle in the highway domain for 10 minutes (Alg. 1, Line 10). We collect their driving data to learn their embedding. This first part of the procedure mirrors the procedure experienced by training participants. All participants next experience four AV conditions as described in Section IV-D (Alg. 1, Line 15). After each condition, participants fill out surveys about their subjective perception of the AV (Section IV-E).

D. Model Testing Study Conditions

The AV behaviors for the four conditions described below are created by shifting a participant’s embedding in the embedding space. Fig. 5 shows the learned embedding space and how we choose the embedding to create the behavior for each of the conditions. We hypothesize that Mimic will produce similar behavior relative to the participant’s driving, Aggressive will produce more aggressive behavior and Cautious, less aggressive. We hypothesize that Perpendicular will produce a driving style that mirrors the aggressiveness of the participant’s driving but modulates other latent aspects of style. By exposing participants to these four conditions, we investigate MAVERIC’s ability to mimic end-user driving style as well as modulate aggressiveness.

Mimic: In *Mimic*, we utilize the personalized embedding learned from the participant’s data to produce driving behavior to mimic the participant’s own driving style.

Aggressive: In *Aggressive*, we shift the participant’s embedding in the positive gradient of aggressiveness (equivalent to 15 points on the ADB survey) to produce more aggressive behavior while maintaining other characteristics of driving style (i.e., $S_\theta(\hat{w}^{(p)}) = \hat{s}^{(p)} + 15$). We constrain $\hat{s}^{(p)} + 15$ to be no more than the largest possible score on the ADB survey (55 points).

Cautious: In *Cautious*, we shift the embedding in the negative gradient of aggressiveness ($S_\theta(\hat{w}^{(p)}) = \hat{s}^{(p)} - 15$) to produce less aggressive behavior while maintaining other characteristics of style. We constrain $\hat{s}^{(p)} - 15$ to be no less than the smallest possible score on the ADB (11 points).

Perpendicular: We include *Perpendicular* to conduct an exploratory investigation into the behavior produced when we maintain the level of aggressiveness but move the embedding within the plane perpendicular to the aggressive gradient. Moving in the plane perpendicular to aggression means that we are able to maintain the level of aggression while modulating other latent aspects of driving style. The plane passing through

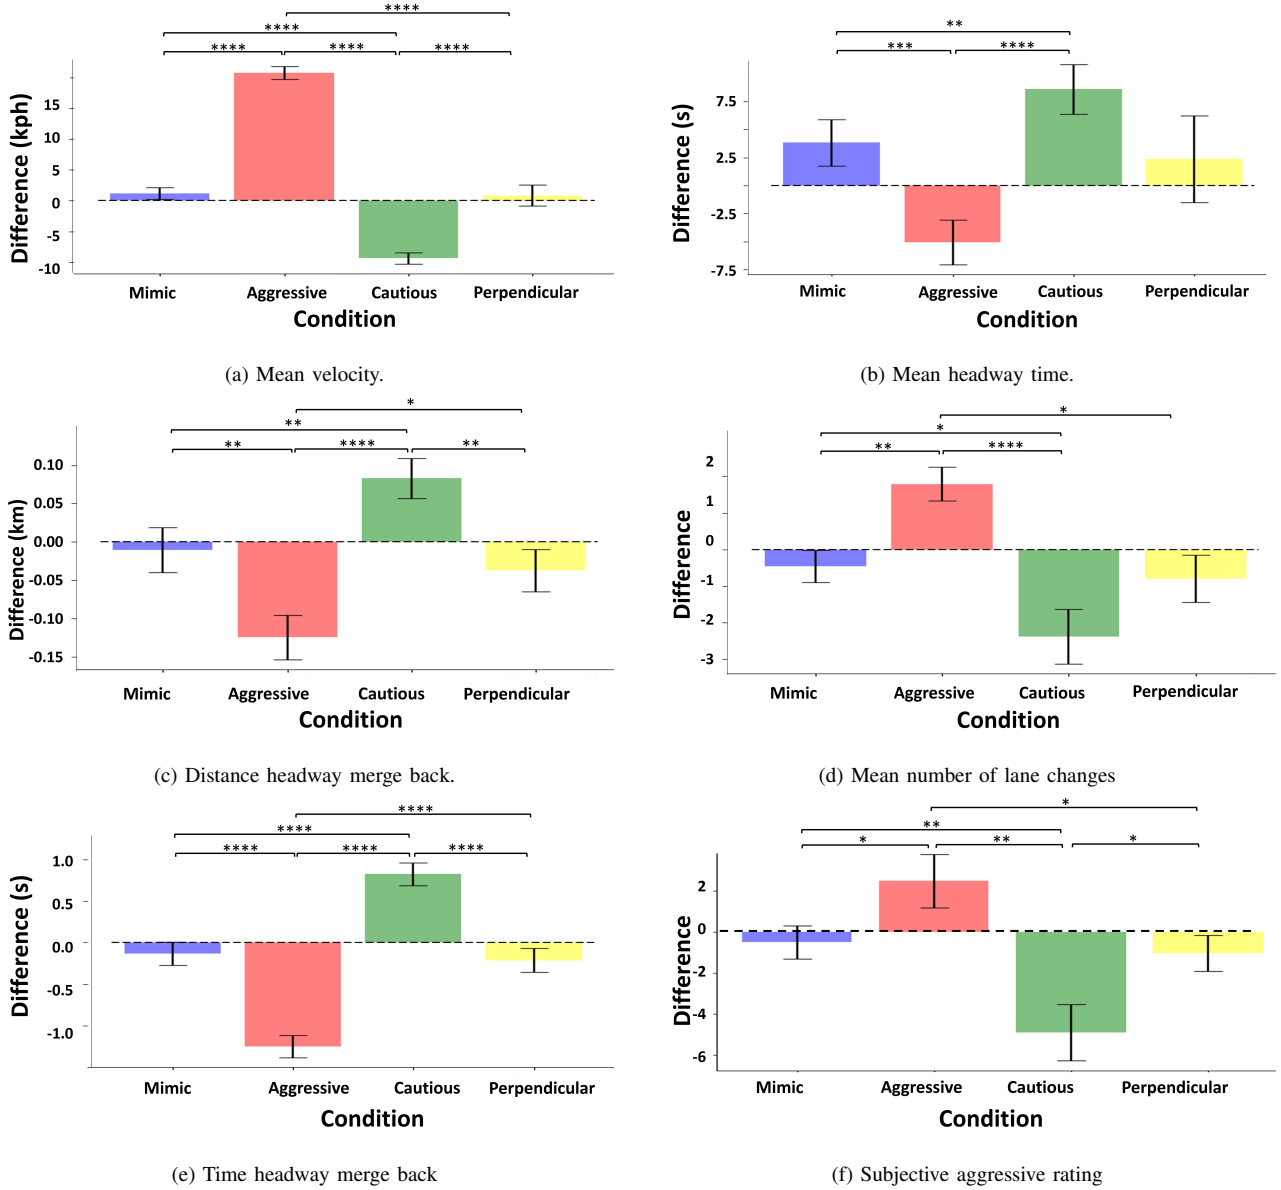


Fig. 6: This figure depicts the difference between the AV’s driving style and the participants’ driving style for our objective and subjective metrics. We show that both objectively and subjectively our approach can mimic an individual’s driving style as well as modulate aggressiveness.

embedding $w^{(p)}$ perpendicular to the aggressive gradient is defined as $\nabla_w S_\theta \cdot (x - w^{(p)}) = 0$ where x is a point in the plane. Our objective is to investigate which driving characteristics change as a result of a shifting $w^{(p)}$ within this plane. To select the embedding, we randomly sample a point along an ellipse on the plane one standard deviation away from the participant’s embedding as shown by the gray points in Fig. 5. By doing so, we are able to keep the degree of aggressiveness constant, while altering other aspects of driver style. We hypothesize that Perpendicular will produce similarly aggressive behavior compared to the participant’s driving but may modulate other factors not related to aggressiveness.

E. Metrics

Participants in both Study 1 and Study 2 complete the pre-study surveys. Only participants in Study 2 complete the

post-trial surveys. The surveys detailed below comply with the design guidelines outlined in Schrum et al. [33] and are validated from prior work when possible.

Pre-study: The pre-study survey is intended to measure the participants’ subjective attitudes towards AVs. We collect demographic information and Big-Five personality information via the Mini International Personality Item Pool [12]. To measure a participant’s aggressive driving style, we utilize the Aggressive Driving Behavior Scale [19]. We measure other aspects of driving style via the Multi-Dimensional Driving Style Inventory [37] and measure experience with cars/racing games/AVs [32], trust in AVs [23], perception of AVs [38], and trust in automation [1].

Post-trial: The post-trial surveys capture the participants’ subjective attitudes towards each of the AV conditions. We measure perceived intelligence [5], competence [9], discomfort

[9], and trust [23]. We modify each of these subscales for AVs. Additionally, we create two custom scales to measure perceived similarity and aggressiveness relative to the participant's own driving style.

Objective Measures: In keeping with prior work [6], [26], we measure various metrics to determine how similar the driving style of each condition is compared to the participant's own driving style. We investigate mean velocity and mean number of lane changes. We also measure mean headway time (the distance between the lead vehicle and ego divided by the speed of the ego when a lane change occurs), minimum headway distance (the minimum distance between the ego and lead vehicle before either the ego slows down or changes lanes), distance headway merge back (the distance between the following vehicle and ego when merging back), and time headway merge back (distance headway merge back divided by the speed when a lane change occurs).

V. RESULTS

A. Analysis of Embedding Space and Aggressive Gradient

We first investigate if our embedding space is capable of representing and producing diverse driving styles and if the aggressive gradient correlates with relevant objective metrics. To investigate these questions, we project the learned embeddings of the test participants onto the line representing the gradient of aggression. We then analyze how driving style changes as a result of the position of the embedding along this line. We find that as we move along the aggressive gradient, the average velocity of the participant increases. The average velocity along the aggressive gradient ranges from 54.5 mph (in the most negative direction of the gradient) to 78.56 mph (in the most positive direction of the aggressive gradient). We find a strong correlation ($r = .49, p = .022$) between the embedding's position along the aggressive gradient and the average velocity of the participant. This finding suggests that, in keeping with prior work [36], velocity is an important component of aggressiveness within the embedding space. We find similar results for mean headway time ($r = -.46, p = .032$), distance headway merge back ($r = -.43, p = .046$), mean number of lane changes ($r = .47, p = .028$), and time headway merge back ($r = -.48, p = .025$). Lastly, we show that a participant's subjective aggressive rating of their own driving style strongly correlates with the position of their learned embedding along the aggressive gradient ($r = .92, p < .001$). **These findings provide evidence that our embedding space is capable of representing diverse driving styles and that aggressiveness objectively and subjectively increases as we move along the aggressive gradient.**

B. Algorithm Validation

We next investigate MAVERIC's ability to mimic end users' driving styles and produce more and less aggressive behavior in terms of both objective and subjective metrics. In our following analysis, we verify that data complies with assumptions before applying a parametric test. We first investigate MAVERIC's ability to accurately mimic driving style. We find that the accuracy with which we are able to mimic

the participant's velocity is 93.6%, time headway is 80.2%, distance headway merge back is 92.4%, mean number of lane changes is 81.0%, and time headway merge back is 81.8%.

Fig. 6 shows the differences in our objective and subjective metrics between the participant's driving style and the behavior produced by our four conditions. To determine if there are significant differences between conditions for each of the metrics, we conduct a repeated measures ANOVA with Holm's post hoc correction or a Friedman's test when the data fails assumptions. We find that the difference between Mimic and the participant's driving is significantly less compared to Aggressive and Cautious for all objective metrics ($p < .001$) (Fig 6a - 6e). We find that Aggressive maintains a higher velocity compared to Mimic. Additionally, as predicted by prior work [6], [26], Aggressive achieves a lower headway merge back time and headway merge back distance. Furthermore, Aggressive commits more lane changes compared to Mimic despite encountering the same number of leading vehicles. We find opposite results with the Cautious condition. We illustrate that the characteristics of our AV driving styles align with the characteristics indicative of aggressiveness in prior work, suggesting that our approach can effectively modulate aggressiveness with respect to one's own driving style [6], [26], [36].




Fig. 7: This figure shows the changes in minimum headway distance as we move around the ellipse within the plane perpendicular to aggressiveness. Minimum headway distance was not significantly correlated with aggressiveness (V-C) and is modulated by moving in the plane perpendicular to aggressiveness.

Additionally, as shown in Fig. 6f, we find that participants rate Cautious as significantly less aggressive compared to Mimic ($p = .002$) and Aggressive as significantly more aggressive ($p = .017$). Furthermore, we find that Mimic and Perpendicular are rated as similarly aggressive compared to the participant's own driving. **Our objective and subjective results together support our hypotheses that 1) our approach is capable of mimicking driving style and 2), by shifting a participant's learned embedding along the**

aggressive dimension, MAVERIC produces objectively and subjectively more aggressive and cautious behavior.

C. Maintaining Other Aspects of Driving Style

One of the goals of our approach is to modulate aggressiveness while maintaining other aspects of driving style. If moving along the gradient of aggressiveness modulates the aggressive aspect of the driving style, then we hypothesize that moving within the plane perpendicular to aggressiveness will modulate other aspects of driving style unrelated to aggressiveness. Interestingly, we found that minimum headway distance and fraction of time in the left lane were not significantly correlated with the embeddings position along the aggressive gradient. Moving along the gradient does not significantly alter minimum headway distance or fraction of time in the left lane, suggesting that, in our learned embedding space, these factors do not play a large role in aggressiveness. Therefore, we predict that these aspects of driving will instead be modulated when we move perpendicular to the gradient of aggressiveness. To test this hypothesis, in Fig 7 we plot the difference in minimum headway distance between Mimic and Perpendicular versus the position around the ellipse that is depicted in Fig. 5. We find that minimum headway distance does in fact correlate with position around the ellipse ($r = .68, p < .001$). We additionally find that the fraction of time in the left lane significantly correlates with position around the ellipse ($r = -.47, p = .025$). Furthermore, Fig 6 shows that participants rate Perpendicular to be the most similar to Mimic across all metrics related to aggressiveness. **Our results support our hypothesis that shifting the personalized embedding within the plane of aggression will maintain aggressiveness while modulating other aspects of driving style.**

We note that minimum headway distance is often associated with aggressiveness [41]. However, this is most often the case when the ego vehicle is not capable of changing lanes and is instead forced to following a leading vehicle. We hypothesize that minimum headway distance is not correlated with aggressiveness because the participant can choose to change lanes at any point to pass a slower driver and therefore is not forced to maintain a following distance if they do not want to.

D. Homophily

As shown in Fig. 8 not all participants preferred the Mimic condition. More than 20% of participants preferred the Aggressive condition and more than 25% of participants preferred the Cautious condition. To explain this finding, we next explore the factors that modulate the effect of homophily (Table I) to determine why some participants prefer a driving style different from their own. First we investigate if a participant's personality impacts their preference via a correlation analysis. As shown in Table I, we find a strong correlation between conscientiousness (i.e., the extent to which one is responsible and dependable [42]) and the difference between a participant's perceived competence of Mimic compared to Cautious, suggesting that individuals higher in conscientiousness prefer a more cautious style to their own. This finding may explain why 62.5% of participants rated

Mimic as less than or equal in competence relative to Cautious. To further support the hypothesis that conscientiousness influences the effect of homophily, we find that participants who are higher in conscientiousness rate a more cautious style as significantly more intelligent, comfortable, and trustworthy compared to Mimic and significantly more competent, intelligent, comfortable, and trustworthy compared to Aggressive.

We additionally find that openness (the degree to which one is broad-minded [42]) correlates with the difference between a participant's comfort with Aggressive compared to Mimic. This finding suggests that those who are more open to new experiences may prefer a more aggressive AV and may explain why 37.5% of participants rated Aggressive as causing greater comfort compared to Mimic.

Prior work suggests that *perceived similarity* to one's own driving style is an important aspect of AV acceptance [6]. To investigate this claim, we conduct a correlation analysis between perceived similarity and an end-users preference for the AV. We find a positive correlation between perceived similarity and trust, intelligence, and competence. This finding suggests that perceived similarity should be taken into consideration when optimizing AV driving style.

Prior work demonstrated that one's own driving style may impact preference for an AV's style (e.g., more aggressive drivers prefer relatively less aggressive AVs) [6], [43]. To investigate this question further, we conduct a correlation analysis between the dimensions of the Multi-Dimensional Driving Style Inventory [37] and preference for Aggressive and Cautious compared to Mimic. We find that participants who report a high-velocity driving style rate Aggressive to be significantly higher than Mimic in intelligence, competence, and trustworthiness. This findings suggests that high-velocity drivers may prefer a more aggressive AV. Due to the contradictory findings with prior work, we aim to conduct a deeper analysis into how the specific dimensions of one's own aggressive style impact the effect of homophily in future work.

We note that the results we present are an exploratory analysis and we do not claim to demonstrate a causal relationship between the subjective factors discussed above and homophily. However, our findings suggest that these factors warrant further investigation in future work. **Overall, our findings demonstrate that personality traits, perceived similarity, and high-velocity driving style may be important factors which influence the effect of homophily.**

VI. DISCUSSION

Our results demonstrate the our MAVERIC framework is capable of both mimicking and modulating driving style by learning an embedding representing an end-user's own driving style. Given other relevant factors related to end-user characteristics, we can then tune this driving style by moving along the gradient of aggression to better match the preference of the end-user. Thus, while other approaches either directly mimic the end-user's own driving style or do not take into consideration the end-user's driving style at all, our approach is capable of integrating both information about an end-user's own driving style and subjective characteristics that are predictive of the optimal AV driving style.

Independent	Dependent	Statistic	p-value
Conscientious	M-C Competence	$\rho(22) = -.71$	$p < .001$
Conscientious	M-C Intelligence	$\rho(22) = -.5$	$p = .012$
Conscientious	M-C Discomfort	$r(22) = .46$	$p = .024$
Conscientious	M-C Trust	$\rho(22) = -.62$	$p = .0011$
Conscientious	M-A Competence	$\rho(22) = -.51$	$p = .011$
Conscientious	M-A Intelligence	$\rho(22) = -.51$	$p = .01$
Conscientious	M-A Discomfort	$r(22) = .45$	$p = .028$
Conscientious	M-A Trust	$\rho(22) = -.48$	$p = .0018$
Openness	M-A Discomfort	$\rho(22) = .49$	$p = .015$
Similarity	Trust	$\rho(94) = .16$	$p = .001$
Similarity	Intelligence	$\rho(94) = .34$	$p < .001$
Similarity	Competence	$\rho(94) = .27$	$p < .001$
High-Velocity	M-A Intelligence	$\rho(94) = -.58$	$p = .0031$
High-Velocity	M-A Competence	$r(22) = -.44$	$p = .03$
High-Velocity	M-A Trust	$r(22) = -.43$	$p = .036$

TABLE I: This table shows our correlation analysis. M represents Mimic, A represents Aggressive, and C represents Cautious.

In our analysis, we show that our learned embedding space captures salient aspects of driving style and that the gradient of aggressiveness correlates with objective and subjective aggressive metrics. An interesting aspect of our aggressive dimension is that this representation of aggressiveness is not based on a pre-defined or hand-crafted heuristic but is instead based upon end-users' perception of what is meant by aggressive driving. By defining aggressiveness via this subjective metric, we are able to produce driving styles that are perceived to be more aggressive or more cautious by end-users rather than relying on an expert's definition which may not align with end-users' perception.

We demonstrate that our approach successfully mimics ($p < .001$) as well as modulates ($p < .001$) driving style with relatively little training data (10 minutes per end-user). Characterizing driving style via a personalized embedding and learning where a driver's style falls within the larger population of drivers improves data-efficiency. While we found that a three-dimensional embedding space was sufficient to represent relevant aspects of driving style in our two-lane highway domain, we hypothesize that increasing the size of the embedding space will enable MAVERIC to capture additional facets of driving style that may appear in more complex domains. We leave to future work an investigation of MAVERIC's performance in these more complex scenarios and the relationship between scenario complexity and the optimal dimensionality of the embedding space.

In our analysis of the effect of homophily, we determine the subjective factors that future work should consider when optimizing driving style. We show that simply mimicking an end-user's own driving style is often not preferred and that certain subjective characteristics may explain the discrepancy between an end-user's own driving style and their preferred AV driving style. By conducting a correlation analysis, we uncover several characteristics that impact homophily. We find that personality ($p < .001$) should be considered when

determining the optimal driving style and that specifically, conscientiousness and openness to experience are important factors. Additionally, participant's perception of their own driving style, e.g. self-reported high-velocity driving style ($p = .0031$) may influence an individual's preference for a more aggressive driving style. We additionally find that perceived similarity ($p < .001$) is a relevant factor as supported in prior work [6]. These findings provide us with insight into what factors should be considered when determining exactly how much and in which direction to shift an end-user's personalized embedding along the aggressive gradient so as to optimize driving style.

VII. FUTURE WORK ON OPTIMIZING DRIVING STYLE VIA MIMICRY AND MODULATION

So far, we have demonstrated that MAVERIC can both mimic and modulate driving style and that several key characteristics of end-users influence the effect of homophily. In future work, we aim to take our approach a step further and characterize the relationship between these relevant subjective metrics and the optimal location of the end-user's personalized embedding along the gradient of aggression. By characterizing this relationship, we can fine-tune the driving style produced by the AV and determine how much and in what direction to shift an end-user's personalized embedding along the gradient of aggression to as to optimize the driving experience for the end-user.

In our current work, we focus on developing a framework that is capable of both mimicking and modulating driving style and leveraging this framework to study factors that influence preference. In future work, we aim to characterize the relationship between these factors and the optimal location of the embedding along the gradient of aggression so as to determine the optimal driving style for an end-user. We then aim to compare this driving style relative to previous approaches. We




Fig. 8: This figure shows the percent of participants who rated Aggressive and Cautious as better than Mimic in terms of each of our subjective metrics.

hypothesize that our framework, which accounts for both an end-user’s own driving style and other relevant latent factors will be able to produce a driving style that is more preferred by end-users compared to prior work.

We chose to modulate the level of aggressiveness because prior work has indicated that this dimension has a large impact on end-user preference [6], [16], [43]. However, we hypothesize that aggressiveness is not the only aspect of driving style that can be modulated with MAVERIC and that other dimensions of driving style may be important to consider. We predict that we can modulate various other aspects of driving style as long as we have access to the appropriate supervision signal when learning the personalized embedding. For example, we could measure an end-user’s preference for environmentally friendly driving and utilize this metric to create an environmentally friendly dimension within embedding space. Additionally, preference for driving style may change over time as the end-user becomes more comfortable with the AV or may change due to circumstance and environmental factors (e.g., inclement weather). In future work, we plan to investigate how to adapt driving style to account for these factors by conditioning the personalized embedding on relevant circumstantial variables.

VIII. LIMITATIONS

A limitation of our work is that we only investigated our approach in a highway domain. In future work, we plan to investigate MAVERIC’s abilities to learn driving styles in domains involving more traffic and the potential for more complex decision making. More complex domains will likely require additional training data to capture the various aspects of driving style. In future work we will investigate training our approach from large offline datasets [10] which we hypothesize will enable MAVERIC to generalize to more complex, real world scenarios.

Additionally, we only recruited internal participants for Study 1. However, despite this limitation, our study comprises

a more diverse population pool than many studies in human-robot interaction which typically recruit from a pool of college students [22]. Another limitation is that the perceived similarity and aggressiveness surveys are not verified in prior work. Because we only conduct a correlation analysis, we cannot conclude that the subjective factors are causally related to homophily, only that they are correlated which is a limitation of our findings. However, our results suggest that these factors are worthy of further investigation in future work and as discussed in VII we aim to conduct a study to quantify the relationship between these factors and the effect of homophily.

IX. CONCLUSION

We have presented MAVERIC, a novel framework for personalizing driving style that uses mimicry as an initial starting point and can fine-tune style by modulating aggressiveness. We demonstrated MAVERIC’s ability to mimic an end-user’s own driving style as well as adjust the level of aggressiveness while maintaining other aspects of style. Furthermore, we employed our framework to study the effect of homophily and showed that preference for one’s own style is modulated by personality, self-reported driving style, and perceived similarity. To our knowledge, ours is the first framework to combine subjective metrics (i.e., the ADB survey) with end-user training data to produce a personalized high-level AV controller. Our results indicate that personalizing AV control is a research area that merits further investigation. We believe that our work provides an important stepping stone towards increasing personalization of AV driving style and ultimately improving end-user experience.

ACKNOWLEDGMENTS

Toyota Research Institute funded this work. This work was supported by the National Science Foundation under grant CNS 2219755.

REFERENCES

- [1] Barbara D Adams, Lora E Bruyn, Sébastien Honde, and Paul Angelopoulos. Trust in Automated Systems Literature Review. *Department of National Defense, Canada*, 2003.
- [2] Il Bae, Jin Hyo Kim, Jaeyoung Moon, and Shiho Kim. Lane change maneuver based on bezier curve providing comfort experience for autonomous vehicle users. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 2272–2277, 2019.
- [3] Il Bae, Jaeyoung Moon, and Shiho Kim. Driving preference metric-aware control for self-driving vehicles. *International Journal of Intelligent Engineering and Systems*, 12:157–166, 2019.
- [4] David Barber and Felix Agakov. The im algorithm: A variational approach to information maximization. In *Proceedings of the 16th International Conference on Neural Information Processing Systems, NIPS’03*, page 201–208, Cambridge, MA, USA, 2003. MIT Press.
- [5] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, pages 71–81, 2009.
- [6] Chandrayee Basu, Qian Yang, David Hungerman, Mukesh Sinahal, and Anca D. Drağan. Do you want your autonomous car to drive like you? In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction*, pages 417–425, 2017.
- [7] Hanna Bellem, Thorben Schönenberg, Josef F. Krems, and Michael Schrauf. Objective metrics of comfort: Developing a driving style for highly automated vehicles. *Transportation Research Part F: Traffic Psychology and Behaviour*, 41:45–54, 2016.
- [8] Drew Bolduc, Longxiang Guo, and Yunyi Jia. Modeling and characterization of driving styles for adaptive cruise control in personalized autonomous vehicles. *Dynamic Systems and Control*, page V001T44A004, 10 2017.
- [9] Colleen M. Carpinella, Alisa B. Wyman, Michael A. Perez, and Steven J. Strossner. The Robotic Social Attributes Scale (RoSAS): Development and Validation. *ACM/IEEE International Conference on Human-Robot Interaction*, Part F1271:254–262, 2017.
- [10] Kan Chen, Runzhou Ge, Heng Qiu, Rami Ai-Rfou, Charles Qi, Xuanyu Zhou, Zoey Yang, Scott Ettinger, Pei Sun, Zhaoqi Leng, Mustafa Mustafa, Ivan Bogun, Weiyue Wang, Mingxing Tan, and Dragomir Anguelov. Womd-lidar: Raw sensor dataset benchmark for motion forecasting, 04 2023.
- [11] Jong Kyu Choi and Yong Gu Ji. Investigating the importance of trust on adopting an autonomous vehicle. *International Journal of Human-Computer Interaction*, 31(10):692–702, 2015.
- [12] Andrew J. Cooper, Luke D. Smillie, and Philip J. Corr. A confirmatory factor analysis of the Mini-IPIP five-factor model personality scale. *Personality and Individual Differences*, 48(5):688–691, 2010.
- [13] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Conference on Robot Learning*, pages 1–16, 2017.
- [14] Mahdi Ebnali, Richard Lamb, Razieh Fathi, and Kevin Hulme. Virtual reality tour for first-time users of highly automated cars: Comparing the effects of virtual environments with different levels of interaction fidelity. *Applied Ergonomics*, 90:103226, 2021.
- [15] Laura Eboli, Gabriella Mazzulla, and Giuseppe Pungillo. How drivers’ characteristics can affect driving style. *Transportation Research Procedia*, 27:945–952, 2017. 20th EURO Working Group on Transportation Meeting, EWGT 2017, 4-6 September 2017, Budapest, Hungary.
- [16] Fredrick Ekman, Mikael Johansson, Lars-Ola Bligård, MariAnne Karlsson, and Helena Strömberg. Exploring automated vehicle driving styles as a source of trust information. *Transportation Research Part F: Traffic Psychology and Behaviour*, 2019.
- [17] Jenny Eriksson and Lars Svensson. Tuning for ride quality in autonomous vehicle application to linear quadratic path planning algorithm. *Thesis, University of Uppsala*, 2015.
- [18] Yingying Feng and Xiaolong Yan. Support vector machine based lane-changing behavior recognition and lateral trajectory prediction. *Computational Intelligence and Neuroscience*, 2022:3632333, May 2022.
- [19] J M Harris and P B Norman. The aggressive driving behavior scale: Developing a self-report measure of unsafe driving practices. *North American Journal of Psychology*, 5:193–202, 2003.
- [20] Martina Hasenjäger and Heiko Wersing. Personalization in advanced driver assistance systems and autonomous vehicles: A review. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–7, 2017.
- [21] M. Hoedemaeker. Driving behaviour with acc and the acceptance by individual drivers. In *IEEE Intelligent Transportation Systems*, pages 506–509, 2000.
- [22] Guy Hoffman and Xuan Zhao. A primer for conducting experiments in human–robot interaction. *J. Hum.-Robot Interact.*, 10(1), 2020.
- [23] Jiun-Yin Jian, Ann Bisantz, and Colin Drury. Foundations for Empirically Determined Scale of Trust in Automated Systems. *International Journal of Cognitive Ergonomics*, 4:53–71, 1998.
- [24] Jesper Karlsson, Sanne van Waveren, Christian Pek, Ilaria Torre, Iolanda Leite, and Jana Tumova. Encoding human driving styles in motion planning for autonomous vehicles. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1050–1056, 2021.
- [25] Markus Kuderer, Shilpa Gulati, and Wolfram Burgard. Learning driving styles for autonomous vehicles from demonstration. *International Conference on Robotics and Automation*, pages 2641–2646, 2015.
- [26] Suzanne E Lee, Erik C B Olsen, Walter W Wierwille, and Michael Goodman. A comprehensive examination of naturalistic lane-changes. *US Department of Transportation Technical Report*, 2004.
- [27] Jiali Ling, Jialong Li, Kenji Tei, and Shinichi Honiden. Towards personalized autonomous driving: An emotion preference style adaptation framework. *IEEE International Conference on Agents*, pages 47–52, 2021.
- [28] Zheng Ma and Yiqi Zhang. Investigating the effects of automated driving styles and driver’s driving styles on driver trust, acceptance, and take over behaviors. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 64(1):2001–2005, 2020.
- [29] Soroush Nasiriany, Huihan Liu, and Yuke Zhu. Augmenting reinforcement learning with behavior primitives for diverse manipulation tasks. *International Conference on Robotics and Automation*, 2021.
- [30] Rohan R. Paleja and Matthew Craig Gombolay. Inferring personalized bayesian embeddings for learning from heterogeneous demonstration. *ArXiv*, abs/1903.06047, 2019.
- [31] Fernando Martín Poó and Ruben Daniel Ledesma. A study on the relationship between personality and driving styles. *Traffic Injury Prevention*, 14(4):346–352, 2013. PMID: 23531257.
- [32] Mariah L. Schrum, Erin Hedlund-Botti, Nina Moorman, and Matthew C. Gombolay. MIND MELD: Personalized Meta-Learning for Robot-Centric Imitation Learning. *ACM/IEEE International Conference on Human-Robot Interaction*, 2022.
- [33] Mariah L. Schrum, Michael Johnson, Muyleng Ghuy, and Matthew C. Gombolay. Four years in review: Statistical practices of likert scales in human-robot interaction studies. *ACM/IEEE International Conference on Human-Robot Interaction*, pages 43–52, 2020.
- [34] Jarrod M Snider. Automatic steering methods for autonomous automobile path tracking. *Thesis: Robotics Institute, Carnegie Mellon University*, 2009.
- [35] Xu Sun, Jingpeng Li, Pinyan Tang, Siyuan Zhou, Xiangjun Peng, Hao Nan Li, and Qingfeng Wang. Exploring personalised autonomous vehicles to influence user trust. *Cognitive Computation*, 12(6):1170–1186, Nov 2020.
- [36] Emilia M. Szumska and Rafał Jurecki. The effect of aggressive driving on vehicle parameters. *Energies*, 13(24), 2020.
- [37] Orit Taubman-Ben-Ari and Vera 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.
- [38] Chris Tennant. *Exploring Emerging Public Attitudes Towards Autonomous Vehicles*, pages 253–266. Springer Singapore, 2021.
- [39] Hanneke Hooft van Huysduynen, Jacques Terken, Jean-Bernard Martens, and Berry Eggen. Measuring driving styles: A validation of the multidimensional driving style inventory. In *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, page 257–264, New York, NY, USA, 2015. Association for Computing Machinery.
- [40] Zhengwei Wu, Paul Schrater, and Xaq Pitkow. Inverse rational control: Inferring what you think from how you forage. *arXiv*, 2019.
- [41] Yueru Xu, Shan Bao, and Anuj K. Pradhan. Modeling drivers’ reaction when being tailgated: A random forests method. *Journal of Safety Research*, 78:28–35, 2021.
- [42] Yangang Yu, Yuanyuan Zhao, Dongyan Li, Jingqiu Zhang, and Jiewei Li. The relationship between big five personality and social well-being of chinese residents: The mediating effect of social support. *Frontiers in Psychology*, 11:1–8, 3 2021.
- [43] Nidzamuddin Md. Yusof, Juffrizal Karjanto, Jacques Terken, Frank Delbressine, Muhammad Zahir Hassan, and Matthias Rauterberg. The exploration of autonomous vehicle driving styles: Preferred longitudinal, lateral, and vertical accelerations. In *Proceedings of the 8th International Conference on Automotive User Interfaces and Interactive Vehicular*

Applications, page 245–252. Association for Computing Machinery, 2016.