

Characterizing Within-Driver Variability in Driving Dynamics During Obstacle Avoidance Maneuvers^{*}

Kendric R. Ortiz^{*} Adam J. Thorpe^{*} AnaMaria Perez^{**}
Maya Luster^{***} Brandon J. Pitts^{***} Meeko Oishi^{*}

^{*} *University of New Mexico, Elec. and Comp. Eng., Albuquerque, NM 87131, USA (e-mail: kendric, ajthor, oishi@unm.edu)*

^{**} *Harvard University, Mathematics, Cambridge, MA 02138, USA (e-mail: anamariaperez@college.harvard.edu)*

^{***} *Purdue University, Industrial Eng., West Lafayette, IN 47907, USA (e-mail: mluster, bjpitts@purdue.edu)*

Abstract: Variability in human response creates non-trivial challenges for modeling and control of human-automation systems. As autonomy becomes pervasive, methods that can accommodate human variability will become paramount, to ensure efficiency, safety, and high levels of performance. We propose an easily computable modeling framework which takes advantage of a metric to assess variability in individual human response in a dynamic task that participants repeat over several trials. Our approach is based in a transformation of observed trajectories to a reproducing kernel Hilbert space, which captures variability in human response as a distribution embedded within the Hilbert space. We evaluate the similarity across responses via the maximum mean discrepancy, which measures the distance between distributions within the Hilbert space. We apply this metric to a difficult driving task designed to elucidate differences across participants. We conducted a pilot study with 6 participants in an advanced driving simulator, in which participants were tasked with collision avoidance of an obstacle in the middle of the road, around a blind corner, in a nighttime scenario, while steering only with the non-dominant hand.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Human-in-the-loop systems, Human modeling, Driving dynamics, Nonparametric learning, Stochastic dynamical systems, Cyber-physical systems.

1. INTRODUCTION

Designing autonomy to be human-centric is paramount for effective and safe human-automation interaction. Indeed, autonomy that is not responsive to the human in the loop can create conditions ripe for “misuse, disuse, [and] abuse” of autonomy (Parasuraman and Riley, 1997). Customization to individual participants is one way to make autonomy human-centric (Cheng et al., 2020; Thropp et al., 2018), and has the potential to significantly impact performance and reliability (de Jonge et al., 2016; Dragan and Srinivasa, 2012) in a manner that can improve human-automation interaction in a wide range of applications.

However, there is a dearth of methods and tools that can effectively characterize heterogeneity across participants, as well as variability of a given participant over time. Historically, models that capture sufficient detail to describe an individual participant have been computationally intractable (Anderson and Lebiere, 1998), or involve non-reproducible modeling processes. Tools that

can capture relevant metrics and can be tailored to an individual, in a repeatable manner, are foundational to the customization and personalization of autonomy for cyber-physical systems (CPS). *In this paper, we propose the use of data-driven techniques to assess variability in individual driver response, based on a metric that captures differences between distributions within a reproducing kernel Hilbert space.* We focus on a driving scenario with naturalistic decision making that confounds traditional control theoretic modeling (Wickens et al., 2021; Weir and McRuer, 1970; Hess, 2000).

In driving assistance systems, customization to driver preferences, characteristics, and behaviors has been shown to be important for system effectiveness (Yi et al., 2020; Hasenjäger and Wersing, 2017). In collision avoidance systems, researchers have investigated data-driven approaches to characterize preferred driver risk (Muehlfeld et al., 2013; Lefevre et al., 2015), and to trigger alerting systems (Wang et al., 2016; Inata et al., 2008) that engage assistive control. Customization impacts controller design in a variety of contexts, and has been investigated in several applications, including fault tolerance (Yi et al., 2020), multi-objective car-following (Gao et al., 2020), and lane-keeping (Rath et al., 2019; Noto et al., 2011), in which the resulting trajectories must be deemed acceptable to

^{*} This material is based upon work supported by the National Science Foundation, under Grant Numbers CNS-1836900 and CNS-1836952. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.



Fig. 1. Participants engaged in a collision avoidance scenario in an advanced driving simulator, steering one-handed, with their non-dominant hand, at night, around an obstacle obscured by a curve in the road.

an individual driver. Other work has focused on reducing potential conflict between the driver and the autonomous system in a shared control framework (Boink et al., 2014) and via potential fields (Noto et al., 2011).

We focus on characterization of within-driver variability in a challenging simulated driving scenario (Figure 1), that is, changes in a single driver's performance over repeated trials, in a scenario carefully constructed to test the participant's ability to adapt to an unfamiliar experience (Luster and Pitts, 2021). This pilot experiment took place in an advanced driving simulator, with a reasonable degree of realism. Gathered data included not only simulated vehicle trajectories, but also eye tracking data.

Our approach to characterize variability is based in a non-parametric machine learning technique, kernel distribution embedding, which exploits properties of reproducing kernel Hilbert spaces. This method enables a data-driven analysis of observed trajectories in a manner that minimizes bias due to tuning, a lack of modeling assumptions, and other heuristics. We project entire trajectories into the Hilbert space, then use the maximum mean discrepancy (MMD) to characterize the distance between the distributions of trajectories in the RKHS. *Our main contribution is a methodological assessment of variability in driving response, via kernel distribution embeddings, and a characterization of its relationship to eye tracking data.*

The organization of the paper is as follows. Section 2 describes the experimental setup. Section 3 provides the mathematical foundation for our proposed approach. Section 4 demonstrates our method on the experimentally gathered data. We discuss implications for control in Section 5 and summarize our findings in Section 6.

2. EXPERIMENTAL SETUP

2.1 Driving scenario

A novel driving scenario was constructed as part of a pilot study at Purdue University, approved under IRB No. 1905022220. The driving scenario consisted of an obstacle avoidance task, in which the participant was to drive in a nighttime setting on a two-lane rural-type highway (Figure 1), and avoid any roadway obstacles that might be present. The driving environment consisted of four S-curve shaped segments, and four straightaways that preceded each S-




Fig. 2. Experimental driving simulator set up, where the obstacle used for avoidance is an old car tire located in the center of the lane at each curve in the road.




Fig. 3. Participants were tasked with avoiding an old tire, which required deviating out of their lane, then returning to their lane as quickly as possible.

curve segment (Figure 2). An obstacle (an old vehicle tire) was placed at the beginning of each S-curve, such that it was not visible to the participant until they rounded the curve. Streetlights were present along both sides of the road for the duration of the drive. There were no other vehicles in the scenario, i.e., no oncoming traffic, and no leading or trailing vehicles.

Participants were asked to drive one-handed, with their non-dominant hand, throughout the entire experiment, to increase the level of difficulty in an otherwise familiar task. Each participant was given the same set of instructions to complete the driving task (Figure 3): (1) practice safe driving (i.e., stay within the lane boundaries), (2) drive at a constant speed of 45 mph throughout the session, and (3) if an obstacle was encountered, avoid the obstacle by moving into the opposite lane as quickly as possible; once cleared, move back into the original lane and resume a constant speed of 45 mph. Participants were first asked to complete a practice drive during daytime on an open highway to familiarize themselves with the driving simulator and task, before engaging in the driving task.

2.2 Collected data and participant demographics

A total of 6 participants completed this pilot experiment. The experiment consisted of three trials, each of which contained four consecutive S-curves; hence, we obtain a total of 12 trajectories for each participant. Trajectories describe position with respect to the centerline of the road (Figure 4), velocity, and heading of the vehicle. Each trial lasted approximately four minutes.

Eye-tracking data was collected as described in Luster and Pitts (2021) for a 7 second duration before each obstacle. An area of interest demarcated an area within the vehicle's lane prior to each obstacle, and average fixation duration




Fig. 4. Each participant generated 12 trajectories. Vehicle position with respect to the centerline is separated for the first 8 trajectories (color) and the last 4 (black).

Participant number	Driving experience
1	6 years
2	3 years
3	1 year
4	6 years
5	2 years
6	7 years

Table 1. Participants' driving experience

was calculated as the average time that the participant was fixated on the area of interest.

Three of the 6 participants were male, and 3 were female; all were senior students in engineering. Average participant age was 21.33 years, with a standard deviation of 0.82 years. Driving experience for each of the participants is summarized in Table 1; the average years of driving experience was 4.17 years.

3. DATA DRIVEN MODELS OF DRIVER-IN-THE-LOOP CPS

3.1 Challenges in modeling human-in-the-loop CPS

Human-in-the-loop systems pose a significant challenge for modeling and control, particularly in configurations in which human response is hard to predict, such as learning scenarios, or other circumstances which invoke naturalistic decision making (Wickens et al., 2021). While a variety of models, of varying degrees of complexity and veracity, have been explored in driving tasks, most are effective under highly constrained scenarios, for which human operators have extensively trained and practiced (Weir and McRuer, 1970; Sheridan, 1992). For the learning scenario we consider, an intentional lack of training and practice confounds these approaches. Hence, one significant challenge is the lack of known mathematical models of the human's inputs to the system.

Presumptions of accurate prior knowledge of the dynamics, as well as of the uncertainty due to the human input, are unrealistic, and could lead to inaccurate assessments, as well as an inability to distinguish features important for

within-driver variability. Further, assumptions regarding the stochastic properties of the system (such as Gaussianity or the Markov property) may simply be inaccurate. Lastly, heterogeneity of human driving behavior means there is likely no single characterization or model which is applicable to all participants, particularly during the learning phase, as drivers gain experience with the driving task.

To sidestep these difficult and potentially unsolvable modeling challenges, we take an empirical, data-driven approach. We model human driver behavior using a statistical learning technique, kernel embeddings of distributions. Although this approach neglects knowledge of the vehicle dynamics, it captures statistical information about the complex interconnections between human inputs and the underlying dynamical system. This approach constructs an implicit representation of the probability distribution of an (unknown) stochastic process in a reproducing kernel Hilbert space (RKHS), which has two key advantages: 1) it is nonparametric, meaning it does not impose prior assumptions on the structure of the uncertainty; 2) it admits a distance metric, the maximum mean discrepancy (MMD), which characterizes the similarity between two distributions.

3.2 Modeling via stochastic processes

We consider the driver-in-the-loop system to be a general stochastic process (Cinlar, 2011),

$$X = \{X_t \mid t \in \mathcal{T}\}, \quad (1)$$

such that for each $t \in \mathcal{T}$, where $\mathcal{T} \in \{0, 1, \dots, N\}$, an index set representing time, for N final time. X_t is a random variable that takes values in the measurable space $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$, where $\mathcal{B}(\mathcal{X})$ is the Borel σ -algebra of $\mathcal{X}$ generated by the set of all open subsets of $\mathcal{X}$. The stochastic process X is a random variable over the space $\Xi = \mathcal{X}^{\mathcal{T}}$, which is the product space of $\mathcal{X}$ over $\mathcal{T}$. We denote the probability distribution of X as $\mathbb{P}$, and realizations of X are called trajectories, denoted $\xi \in \Xi$.

Using data collected from the driving experiment, we seek to compute an empirical representation of the distribution of X (1) as an element in an RKHS.

Definition 1. (RKHS, Aronszajn, 1950). A Hilbert space $\mathcal{H}$ of functions from Ξ to $\mathbb{R}$ is a reproducing kernel Hilbert space (RKHS) if there exists a positive definite function $k : \Xi \times \Xi \rightarrow \mathbb{R}$, called the kernel, that satisfies the following properties:

- (1) For every $\xi \in \Xi$, $k(\xi, \cdot) \in \mathcal{H}$, and
- (2) (Reproducing property): For every $\xi \in \Xi$ and $f \in \mathcal{H}$, $f(\xi) = \langle f, k(\xi, \cdot) \rangle_{\mathcal{H}}$.

A common choice of kernel function is the Gaussian kernel, $k(\xi, \xi') = \exp(-\|\xi - \xi'\|^2 / 2\sigma^2)$. Furthermore, the RKHS is equipped with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and the implicitly defined norm $\|\cdot\|_{\mathcal{H}}$.

At the heart of our approach is an element $m_{\mathbb{P}} \in \mathcal{H}$, the kernel distribution embedding of $\mathbb{P}$ (Smola et al., 2007), defined as

$$m_{\mathbb{P}} = \int_{\Xi} k(\xi, \cdot) \mathbb{P}(d\xi), \quad (2)$$

where $\mathbb{E}[f(X)] = \langle f, m_{\mathbb{P}} \rangle_{\mathcal{H}}$ by the reproducing property. This representation allows us to capture the probability distribution of X as an element of the RKHS.

However, in practice, we cannot compute (2) directly, since the distribution $\mathbb{P}$ of X is unknown. Hence, following Smola et al. (2007), we compute an empirical estimate $\widehat{m}_{\mathbb{P}}$ from data $\mathcal{D} = \{\xi^i\}_{i=1}^M$, collected i.i.d. from $\mathbb{P}$, such that

$$\widehat{m}_{\mathbb{P}} = \frac{1}{M} \sum_{i=1}^M k(\xi^i, \cdot). \quad (3)$$

In effect, this enables empirical estimation of the probability distribution of the driver-in-the-loop system as an element of the RKHS.

3.3 Maximum mean discrepancy

One particularly useful property of the RKHS facilitates analysis of the similarity of functions and distributions. The implicit distance metric in the RKHS $\mathcal{H}$, $\|f - g\|_{\mathcal{H}}$, where $\|f - g\|_{\mathcal{H}} = 0$ implies $f = g$, can be used to measure the similarity of two distributions $\mathbb{P}, \mathbb{Q}$ on Ξ via their embeddings, $m_{\mathbb{P}}$ and $m_{\mathbb{Q}}$,

$$\text{MMD}(\mathbb{P}, \mathbb{Q}) = \|m_{\mathbb{P}} - m_{\mathbb{Q}}\|_{\mathcal{H}}. \quad (4)$$

This is the *maximum mean discrepancy* (MMD) (Gretton et al., 2012).

Importantly, if the kernel is *characteristic* (Sriperumbudur et al., 2010), then the embedding is unique, meaning it captures all statistical features of the underlying distribution and is injective (Smola et al., 2007, Theorem 1). In other words, if k is characteristic, then for any two distributions $\mathbb{P}$ and $\mathbb{Q}$, $\|m_{\mathbb{P}} - m_{\mathbb{Q}}\|_{\mathcal{H}} = 0$ if and only if $\mathbb{P} = \mathbb{Q}$. Many common choices of kernels are characteristic, including the Gaussian kernel that we employ.

Using the empirical estimate (3), we can compute an estimate of (4) as

$$\widehat{\text{MMD}}(\mathbb{P}, \mathbb{Q}) = \left\| \frac{1}{M} \sum_{i=1}^M k(\xi^i, \cdot) - \frac{1}{N} \sum_{i=1}^N k(z^i, \cdot) \right\|_{\mathcal{H}}. \quad (5)$$

Thus, we can readily employ the MMD to assess the similarity between empirical estimates of distributions.

We note that although the maximum mean discrepancy is a type of integral probability metric (Sriperumbudur et al., 2012) (such as the KL divergence, or total variation divergence), it does not require density estimation, which can be computationally expensive (Sriperumbudur et al., 2010; Gretton et al., 2012). This computational ease makes the MMD a good candidate for near-run time applications for the purpose of control. Additionally, alternative approaches, such as those based in moment matching, may not only be difficult to employ with limited data, but also computationally expensive, as a high number of moments may be required to ensure sufficient accuracy.

4. ANALYSIS AND RESULTS

4.1 Computing empirical estimates of embeddings

We presume the observed data is of the form

$$\mathcal{D} = \{\xi^i\}_{i=1}^M \quad (6)$$

where $M \in \mathbb{N}$ denotes the total number of trajectories in the data set $\mathcal{D}$, and each trajectory $\xi^i \in \mathcal{D}$ is a sequence of state vectors indexed by time,

$$\xi^i = \{x_0^i, x_1^i, \dots, x_N^i\} \quad (7)$$

where $x_t^i \in \mathcal{X}$ for all $i = 1, \dots, M$ and $t = 0, 1, \dots, N$. We denote $\mathcal{D}_j$ to be the set of all trajectories in $\mathcal{D}$ corresponding to the j^{th} participant. We assume that the trajectories of each participant j are drawn i.i.d. from stochastic process X^j with unknown probability distribution $\mathbb{P}_j$, where the stochasticity is due to the human-in-the-loop operator.

Hence for each participant j , with $M = 12$, we have data $\mathcal{D}_j = \{\xi_j^1, \xi_j^2, \dots, \xi_j^{12}\}$, with trajectories $\xi_j^i \in \mathbb{R}^{3N}$ over $N = 1100$ time steps, at a sampling rate of 60 Hz.

We first choose to separate the first eight trajectories ($\xi_j^1, \dots, \xi_j^8$) from the last four trajectories ($\xi_j^9, \dots, \xi_j^{12}$), which describe the *final distribution*. We seek to compare each of the first eight trajectories with the final distribution to assess consistency of driving performance over time.

This separation of the last four trajectories for comparison with each of the first 8 trajectories was chosen largely to exploit the fact that the experimental setup is intentionally designed for participants to increase their familiarity with the driving task as the trials progress. By the ninth trajectory, we expect that the participant has become quite familiar with the trajectory, as supported by recognition-primed decision making (RPDM) (Fineberg et al., 2000; Ross et al., 2004). The theory of RPDM indicates that recognition based strategies, as opposed to analytically driven strategies, are employed when familiar scenarios are encountered. Although other separation strategies could have been chosen, we choose to use the natural grouping suggested by the experiment, in which the last four trajectories were completed as a unit (trial 3).

We compute the empirical kernel distribution embedding of the final distribution for each j^{th} participant from (3), as

$$\widehat{m}_{\mathbb{P}_j} = \frac{1}{4} \sum_{i=9}^{12} k(\xi_j^i, \cdot) \quad (8)$$

for $\xi_j^i \in \mathcal{D}_j$. We used a Gaussian kernel with a fixed bandwidth parameter $\sigma = 32$, chosen via cross validation to ensure that the MMD can distinguish between distributions.

We then compute the j^{th} participant's maximum mean discrepancy for each of the first 8 trajectories with respect to (8), by evaluating (5) for each, as

$$\widehat{\text{MMD}}(\xi_j^i, \mathbb{P}_j) = \|k(\xi_j^i, \cdot) - \widehat{m}_{\mathbb{P}_j}\|_{\mathcal{H}}. \quad (9)$$

Since the Gaussian kernel is characteristic, the MMD will decrease as trajectories become more consistent with the final distribution, and the MMD will increase when trajectories became more dissimilar.

4.2 Trajectory analysis via maximum mean discrepancy

As shown in Figure 5, there is considerable heterogeneity in trends across participants. Each MMD value indicates




Fig. 5. Maximum mean discrepancy for all 6 participants in as a function of time. Participants 1, 3, and 4 show decreasing trends; participants 5 and 6 show variability but no overall change; participant 2 shows an initial decrease over the first four trajectories, followed by an increase over the next four trajectories.

the respective dissimilarity between the compared distributions. For example, for participant 1, we can see that the the largest MMD value is approximately 1.25 for the 2nd trajectory, and the lowest MMD value is approximately 0.7 for the seventh trajectory. These numbers indicate that the second trajectory is the most dissimilar from the participant's *final distribution*, and the seventh trajectory is the most similar.

Participants 1, 3, and 4 show a decreasing trend in MMD as trajectories progress, indicating that their performance became more consistent with the final distribution as the participants gained familiarity with the obstacle avoidance task. We interpret this trend as evidence of learning, due to the improvements in performance over time.

Participants 5 and 6 show no major increases or decreases in MMD over the duration of the entire experiment. Participant 5 has a fairly consistent MMD, whereas Participant 6 shows considerable variability, with MMD values that range between 0.9 and 1.1. In both cases, we infer that driving performance was essentially unchanged in its variability over the duration of the first two trials (i.e., first 8 trajectories) of the experiment. There are several potential explanations for this finding, including a ceiling effect (i.e., maximum performance has been reached with little to no room for improvement), a lack of environmental stimuli (e.g., no on-coming, opposing traffic) to encourage improvement, no motivation for enhancing performance given the absence of experiment incentives, and/or behavior being driven by one's own expectation and interpretation of the goal(s) of the study.

Participant 2 shows a completely different trend: this participant has a decrease in MMD over the first four trajectories, but an increase in MMD over the next four trajectories. We believe that this is due to two primary factors: 1) a loss of attention that was informally observed during the experiment, and 2) a poor understanding of the instructions, evident, in part, by the participant's lack of




Fig. 6. Average fixation duration in the 7 seconds before the obstacle, for the first 8 trajectories for all 6 participants.

focus as the experiment progressed, as well as the fact that the participant employed two hands (contrary to instructions) for the first four trajectories.

4.3 Correlation with eye-tracking analysis

Since eye tracking analysis can provide insights into attention allocation, we also consider the relationship between driving performance (indicated by MMD) and eye behavior, i.e., average fixation duration, a measure of the amount of time a person visually engages with particular area of interest. We consider the average fixation duration in the 7 seconds before each obstacle for all 6 participants (Figure 6), and construct a linear regression with respect to MMD (Figures 7 and 8).

Participants that exhibit learning (i.e., those who have improvements in driving performance as the experiment progresses) have divergent trends (Figure 7a). This possibly indicates different strategies employed by participants. For example, Participant 1 showed a decrease in attention over time, which is consistent with the hypothesis in (Luster and Pitts, 2021), that as learning occurs, a shorter fixation duration is needed to complete an identical task, given that familiarity is built based on previous trials. In contrast, Participant 4 shows increased attention over time, possibly due to their belief that more attention was required to achieve improved performance. We believe that the difference in trends between Participant 3 and Participant 4 may be evidence of different learning strategies. The former may show less effort required to achieve improved performance, whereas the latter may be evidence of more effort at participants' discretion.

Participants 5 and 6, who do not exhibit clear trends in learning, show relatively static linear fit trends in eye-tracking, as well. In essence, both driving performance and attention appear to be unchanged for the duration of the experiment. We note that while Participant 3 also has a relatively constant eye-tracking linear fit (see Figure 7a), their performance is more consistent with Participants 1 and 4, than it is with Participants 5 and 6 (as shown in Figure 5). This suggests that Participant 3 might have not




Fig. 7. The relationship between maximum mean discrepancy and average fixation duration shows different trends for participants who (a) exhibited learning, and (b), those who did not. Linear regression indicates small absolute slopes for those participants who did not show improvements in MMD, and larger absolute values of slopes for those who did.

adopted any particular attentional strategies in successive trials, which by chance, resulted in marginally poorer driving performance over time. However, confirmation of this hypothesis would require a larger data set with additional participants beyond this pilot study, and partitioning them into groups based on distinctive similarities.

Lastly, Participant 2 shows a trend in average fixation duration that is similar to that of Participants 5 and 6, and is consistent with the lack of attention informally observed for this participant. We note that the linear regression in Figure 8a may be difficult to interpret because it involves data from essentially two different stochastic processes, that is: two-handed driving in the first four trajectories (trial 1), and one-handed driving in the next four trajectories (trial 2). Separating the data into the two separate trials Figure 8b results in steep regression slope for the first four trajectories (comparable to that of Participant 4), and a shallow regression slope for the next four trajectories (comparable to that of Participants 5 and 6). We believe this shows evidence of learning in the first trial (first four trajectories), and disengagement in the second trial.

5. IMPLICATIONS FOR CONTROL

One of the major contributions of this work is the quantitative assessment of 1) within-participant variability in a difficult driving task, as well as 2) heterogeneity across participants in those trends. The variability and heterogeneity demonstrated here has potential for significant impact in control and in autonomous systems. Specifically, a one-size-fits all control could have a detrimental impact on some participants. Consider a controller designed to help improve performance of novices, by helping them learn the desired task more quickly than they might learn it on their own. For example, a controller designed to intervene in response to increasing fixation durations could actually impair the learning process for participants whose learning strategies involve increased attention at their performance improves. Similarly problematic phenomena could arise in




Fig. 8. (a) Maximum mean discrepancy and average fixation duration for Participant 2 during the the first eight trajectories, with linear regression. (b) The same data, separated into trial 1 (i.e., the first four trajectories, marked with '+') and trial 2 (i.e., the fifth through eighth trajectories, marked with '*'). Linear regressions have vastly different slopes, indicating learning during trial 1, but not during trial 2.

driver assistance systems designed to detect exceptions to anticipated phenomena, for example.

We maintain that development of future autonomous systems will require accurate assessments of individual performance, customized not only to the individual, but also responsive to the operational context of human CPS.

6. CONCLUSIONS

We provide a data-driven approach to assess variability in a difficult driving scenario. Because mathematical models are scarce for naturalistic scenarios such as those that involve learning, our approach exploits kernel distribution

embeddings, which limit potential inaccuracies posed by incorrect model assumptions. We characterize variability in the driving task, and show potentially different learning strategies for those participants that show evidence of learning through driving performance. Based on our preliminary work here, we propose that that characterization of heterogeneity is indeed important for design of autonomy that is customizable and truly human-centric, and that accounting for individual variability is important for future work in customizable control of human CPS.

REFERENCES

- Anderson, J. and Lebiere, C. (1998). *The Atomic Components of Thought*. Lawrence Erlbaum Assoc., NJ.
- Aronszajn, N. (1950). Theory of reproducing kernels. *Trans. of the Amer. Math. Soc.*, 68(3), 337–404.
- Boink, R., van Paassen, M., Mulder, M., and Abbink, D. (2014). Understanding and reducing conflicts between driver and haptic shared control. In *IEEE Int'l Conf. on Syst., Man, Cybern.*, 1510–1515.
- Cheng, Y., Sun, L., Liu, C., and Tomizuka, M. (2020). Towards efficient human-robot collaboration with robust plan recognition and trajectory prediction. *IEEE Robot. Autom. Lett.*, 5(2), 2602–2609.
- Cinlar, E. (2011). *Probability and Stochastics*. Springer.
- de Jonge, A., Wildenbeest, J., Boessenkool, H., and Abbink, D. (2016). The effect of trial-by-trial adaptation on conflicts in haptic shared control for free-air teleoperation tasks. *IEEE Trans. Haptics*, 9(1), 111–120.
- Dragan, A. and Srinivasa, S. (2012). Online customization of teleoperation interfaces. In *IEEE Int'l Symp. on Robot and Human Interactive Commun.*, 919–924.
- Fineberg, M., Boff, K., Chapanis, A., Rouse, W., and OH., C.S.E.I.A.C.W.P.A. (2000). *Gateway: Volume 11, Number 3, 2000*. Defense Technical Information Center. URL <https://books.google.com/books?id=zv0FngAACAIAJ>.
- Gao, B., Cai, K., Qu, T., Hu, Y., and Chen, H. (2020). Personalized adaptive cruise control based on online driving style recognition technology and model predictive control. *IEEE Trans. Veh. Technol.*, 69(11), 12482–12496.
- Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., and Smola, A. (2012). A kernel two-sample test. *J. Mach. Learn. Res.*, 13(25), 723–773.
- Hasenjaeger, M. and Wersing, H. (2017). Personalization in advanced driver assistance systems and autonomous vehicles: A review. In *IEEE Int'l Conf. on Intell. Transp. Syst.*, 1–7.
- Hess, R. (2000). Human-in-the-loop control. In *Control System Applications*, 327–334. CRC Press Inc.
- Inata, K., Raksincharoensak, P., and Nagai, M. (2008). Driver behavior modeling based on database of personal mobility driving in urban area. In *Int'l Conf. on Ctrl., Autom. and Syst.*, 2902–2907.
- Lefevre, S., Carvalho, A., Gao, Y., Tseng, E., and Borrelli, F. (2015). Driver models for personalized driving assistance. *Veh. Syst. Dyn.*, 53.
- Luster, M. and Pitts, B. (2021). A preliminary investigation into learning behaviors in complex environments for human-in-the-loop cyber-physical systems. In *Human Fac. and Ergo. Soc.*, volume 65, 42–46.
- Muehlfeld, F., Doric, I., Ertlmeier, R., and Brandmeier, T. (2013). Statistical behavior modeling for driver-adaptive precrash systems. *IEEE Trans. Intell. Transp. Syst.*, 14(4), 1764–1772.
- Noto, N., Okuda, H., Tazaki, Y., Inagaki, S., and Suzuki, T. (2011). Obstacle avoidance assisting system based on personalized potential field. In *SICE Annual Conf. 2011*, 476–481.
- Parasuraman, R. and Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Fac.*, 39(2), 230–253.
- Rath, J., Senouth, C., and Popieul, C. (2019). Personalised lane keeping assist strategy: adaptation to driving style. *IET Ctrl. Theory & Appl.*, 13(1), 106–115.
- Ross, K., Klein, G., Thunolm, P., Schmitt, J., and Baxter, H. (2004). The recognition-primed decision model. *Military review*, 6–10.
- Sheridan, T. (1992). *Telerobotics, automation, and human supervisory control*. MIT press.
- Smola, A., Gretton, A., Song, L., and Schölkopf, B. (2007). A hilbert space embedding for distributions. In *Int'l Conf. Alg. Learning Theory*, 13–31. Springer.
- Sriperumbudur, B., Fukumizu, K., Gretton, A., Schölkopf, B., and Lanckriet, G. (2012). On the empirical estimation of integral probability metrics. *Elec. J. of Stat.*, 6, 1550 – 1599.
- Sriperumbudur, B., Gretton, A., Fukumizu, K., Schölkopf, B., and Lanckriet, G. (2010). Hilbert space embeddings and metrics on probability measures. *J. Mach. Learn. Res.*, 11, 1517–1561.
- Thropp, J., Oron-Gilad, T., Szalma, J., and Hancock, P. (2018). Calibrating adaptable automation to individuals. *IEEE Trans. Human-Mach. Syst.*, 48(6), 691–701.
- Wang, J., Yu, C., Li, S., and Wang, L. (2016). A forward collision warning algorithm with adaptation to driver behaviors. *IEEE Trans. Intell. Transp. Syst.*, 17(4), 1157–1167.
- Weir, D. and McRuer, D. (1970). Dynamics of driver vehicle steering control. *Automatica*, 6(1), 87–98.
- Wickens, C., Helton, W., Hollands, J., and Banbury, S. (2021). *Engineering psychology and human performance*. Routledge.
- Yi, D., Su, J., Hu, L., Liu, C., Quddus, M., Dianati, M., and Chen, W.H. (2020). Implicit personalization in driving assistance: State-of-the-art and open issues. *IEEE Trans. Intell. Veh.*, 5(3), 397–413.