

Toward Informed AV Decision-Making: Computational Model of Well-being and Trust in Mobility

Zahra Zahedi, Shashank Mehrotra, Teruhisa Misu and Kumar Akash

Honda Research Institute USA, Inc

{zahra_zahedi, shashank_mehrotra, tmsu, kakash}@honda-ri.com

Abstract

For future human-autonomous vehicle (AV) interactions to be effective and smooth, human-aware systems that analyze and align human needs with automation decisions are essential. Achieving this requires systems that account for human cognitive states. We present a novel computational model in the form of a Dynamic Bayesian Network (DBN) that infers the cognitive states of both AV users and other road users, integrating this information into the AV’s decision-making process. Specifically, our model captures the “well-being” of both an AV user and an interacting road user as cognitive states alongside trust. Our DBN models infer beliefs over the AV user’s evolving well-being, trust, and intention states, as well as the possible well-being of other road users, based on observed interaction experiences. Using data collected from an interaction study, we refine the model parameters and empirically assess its performance. Finally, we extend our model into a causal inference model (CIM) framework for AV decision-making, enabling the AV to enhance user well-being and trust while balancing these factors with its own operational costs and the well-being of interacting road users. Our evaluation demonstrates the model’s effectiveness in accurately predicting user’s states and guiding informed, human-centered AV decisions.

1 Introduction

With the proliferation of autonomous vehicles (AVs), including cars and even smaller autonomous vehicles such as delivery robots and drones, humans are bound to encounter AVs in more places and in different forms. However, existing research on automation techniques overlooks the environmental and social implications while developing these systems [Zhuge and Zhuge, 2020]. For example, recently, some of these new mobility modes that had been considered convenient and climate-friendly have received a public perception of being dangerous and a nuisance. For example, Paris referendum, 89% of voters supported a ban on electric scooters [Nouvian, 2023]. One of the primary reasons for such public sentiment is that these shared mobility modes only con-

sider the basic mobility needs of the users with disregard for their well-being, satisfaction, and positive relationship with other road users [Ettema *et al.*, 2011]. Given the recent convergence between automated vehicle (AV) technology and shared mobility, new small-scale shared automated vehicle tests are beginning to develop around the world [Stocker and Shaheen, 2017]. Therefore, it becomes even more critical for these AVs to holistically consider the needs of the users as well as nearby road users. A potential paradigm to achieve this is to account for human cognitive states while making AV decisions. Studies have established a strong association between driving behavior and well-being [Harris *et al.*, 2014]. Similarly, a user’s trust influences how users perceive the AV’s intentions and reliability, which in turn affects their willingness to accept its decisions [Lee and See, 2004]. Thus, it is necessary for the AV to quantitatively measure users’ well-being and trust to better anticipate the impact of its actions on both its user and others on the road. We propose to focus on *well-being* as a cognitive state to account for holistic user needs that include situational satisfaction as well as positive relation with other road users, and to ensure successful interaction between the user and their AV, we incorporate user’s *trust* in the AV.

In this paper, we propose a Bayesian-based model that allows for the inference of well-being and trust without disrupting the user’s behavior. The proposed model can continuously maintain an estimate of the user’s cognitive states and update it based on the user’s and other road users’ latent states. We use Lee and See’s definition of user’s trust as their attitude toward AV/others that they will help them achieve their goal of driving safely in a situation characterized by uncertainty and vulnerability [Lee and See, 2004]. Well-being is a multidimensional construct and constitutes several factors [Seligman and others, 2002]. For the mobility context, we define well-being to comprise of positive social interactions, satisfaction with travel, trust in other road users, and general well-being of the user [Radzyk, 2014]. Furthermore, our work introduces the novel application of incorporating inferred well-being into decision-making processes, a unique contribution that has not been explored in previous research. The main contributions of the paper include:

1. Development of a quantitative computational model of trust and well-being based on a dynamic Bayesian network (DBN) that can infer user’s trust, well-being, and

intention. The model was then trained and evaluated using the collected user study data.

2. Inference of optimal decision-making policies through interaction to achieve desired objectives, such as maximizing the user’s well-being and trust as well as optimizing a trade-off between the well-being of the user, other road users, and AV costs.

2 Related Work

The consideration of cognitive states in human-automation interactions has been an emerging research focus, particularly as autonomous systems become more integrated into everyday environments. Several studies have examined the importance of understanding cognitive states, such as trust, attention, workload, and situational awareness to ensure effective human-automation collaboration [Akash *et al.*, 2020; Azevedo-Sa *et al.*, 2020; Wu *et al.*, 2022]. These studies emphasize the need for systems that can interpret and respond to these states, with the aim of enhancing user experience and safety.

One key challenge is the dynamic nature of cognitive states, which evolve over time based on the interaction context. Bayesian inference and modeling have been widely employed in various domains for modeling cognitive states, including workload [Koppol *et al.*, 2021; Luo *et al.*, 2019; Guhe *et al.*, 2005], trust [Xu and Dudek, 2015; Guo *et al.*, 2020; Soh *et al.*, 2020], distraction [Liang *et al.*, 2007; Liang and Lee, 2014; Zahedi *et al.*, 2022], and emotion [Ong *et al.*, 2019]. These approaches have yielded valuable insights into human behavior and decision-making processes in diverse domains [Mahmood *et al.*, 2024; Zahedi *et al.*, 2023; Xu and Dudek, 2016; Luo *et al.*, 2021; Deo and Trivedi, 2019].

In addition to the computational modeling of cognitive states, the relationship between driving and well-being has been an area of research. Various studies have explored how factors like stress, fatigue, and emotional states affect driver behavior and safety, providing indirect evidence of the connection between cognitive states and well-being. For example, levels of well-being are correlated with driving performance [Hu *et al.*, 2013; Bowen and Smith, 2019]. and levels of driving violations [Isler and Newland, 2017]. Also, prosocial driving behavior promotes cooperation with other road users and reduces incidents of aggressive and stressful driving [Harris *et al.*, 2014]. While these studies contribute valuable insights, they typically rely on self-reported measures or physiological indicators to assess well-being. Self-report questionnaires, such as [Radzyk, 2014; Friman *et al.*, 2013] and physiological measurements, such as [Sauer *et al.*, 2019; Halkola *et al.*, 2019] are commonly used to assess subjective states of well-being. However, these measures are potentially distracting or intrusive; they may not be practical for real-time decision-making in safety-critical environments. Mehrotra *et al.* explored the factors impacting well-being and trust and proposed a support vector machine model to understand these factors [Mehrotra *et al.*, 2023]. However, to create a comprehensive model, we require an informed model

that incorporates cognitive structures and accounts for the dynamic nature of well-being and trust.

3 Problem Formulation

Our model is predicated on three key relationships: (1) relationship between the user (denoted as E) and their AV (i.e. R), (2) relationship between the AV and the other road user (denoted as O), and (3) relationship between the user and the other road user. We consider a dyadic bi-directional interaction that involves possible symmetric actions. At each interaction, either the other road user or AV can contribute toward the other by an accommodative action, and the other is the receiver of that. Accommodative action can be choosing a positive prosocial action toward others (R_+ for AVs as action contributors, or O_+ for others contributing to the action) or not (R_-/O_-). When the AV is the contributor, the user of the AV might have an intention for accommodative action toward the other (I_+ as positive intention toward accommodative action or not I_-) that may or may not align with the one the AV chooses (Al_1 or Al_0 respectively). Therefore, users well-being, trust, and their action-alignment can affect the relationship between the user and the AV.

Formally, the goal of this work is to infer the degree of user’s wellbeing $w_k \in [0, 1]$ (where 0 is the lowest and 1 is the highest), trust on the AV $t_k \in [0, 1]$ and the intention towards others $i_k \in \{I_+, I_-\}$, as well as other’s wellbeing $w_k^O \in [0, 1]$, at each interaction event $k = 1 : K$. We tackle this problem by relating the latent states to observable factors of AV’s and others accommodative actions ($a_k^R \in \{R_+, R_-\}$ and $a_k^O \in \{O_+, O_-\}$ respectively). Since at any interaction event, the user intention might not necessarily be the same as the AV action; we consider an action-alignment state, represented by $al_k \in \{Al_0, Al_1\}$ where Al_0 indicates not-aligned and Al_1 indicates aligned with the user’s intention.

3.1 Interaction Models

We formalize the interaction model as two Dynamic Bayesian Networks (DBN). We could have modeled the interaction as a 2-step Temporal Bayesian Network (2TBN), in which the O and R alternatively contribute an accommodative action. However, in order not to limit the framework to alternative interaction between others and AV as an action contributor, we modeled each interaction as a separate DBN, one for when the AV is taking an action (R-DBN, or R contributor DBN), and the other for when the other is taking an action (O-DBN, or O contributor DBN).

At each event k , our model treats the state of user’s latent states and others’ states as random variables and maintains belief distributions based on various factors of the interaction experience. Our R-DBN model relates the user’s latent states $x_k^E = \{w_k, t_k, i_k\}$ causally to the AV’s action a_k^R , other’s action a_k^O and user’s intention alignment state with the AV al_k , and The O-DBN relates the other’s latent state $x_k^O = w_k^O$ to AV’s action a_k^R , their action a_k^O . Links to each factor in these Bayesian models are quantified as a conditional probability distribution (CPD). How the user’s states x_k^E expected to change given the current AV’s action a_k^R and action alignment al_k and possible recent others’ actions a_{k-1}^O is reflected




Figure 1: R-DBN (blue) and O-DBN (red) in general form in alternative order.

as follows when AV is action contributor and when other is action contributor respectively:

$$x_k^E \sim \text{Prob}(x_k^E | x_{k-1}^E, a_k^R, al_k) \quad (1)$$

and

$$x_k^E \sim \text{Prob}(x_k^E | x_{k-1}^E, a_{k-1}^O) \quad (2)$$

Similarly, for other's state x_k^O the expected change when the AV is contributing to the action is represented as

$$x_k^O \sim \text{Prob}(x_k^O | x_{k-1}^O, a_k^R) \quad (3)$$

When others are contributing to an action as

$$x_k^O \sim \text{Prob}(x_k^O | x_{k-1}^O, a_k^O) \quad (4)$$

This probabilistic representation allows us to infer the expected human states and other's state at any given event, as well as the level of uncertainty associated with each estimate. The general graphical model of the casual and evidential variable interaction are shown in Figure 1.

3.2 Bayesian Inference and Prediction

The models allow us to estimate the probabilistic belief over the user's latent states x_k^E at event k . Inference is performed by computing the posterior distribution of x_k^E or x_k^O given past observations. Using Bayesian filtering, we recursively compute

$$P(x_k^E | \text{evidence}_{1:k}) \propto P(\text{evidence}_k | x_k^E) \times \sum_{x_{k-1}^E} P(x_k^E | x_{k-1}^E, a_k^R, al_k, a_{k-1}^O) P(x_{k-1}^E | \text{evidence}_{1:k-1}) \quad (5)$$

where evidence_k consists of observed variables such as the AV's action a_k^R , others' actions a_k^O , and the alignment state al_k .

For prediction, given an initial belief $P(x_0^E)$, the future state can be estimated by marginalizing over latent states:

$$P(x_{k+1}^E | \text{evidence}_{1:k}) = \sum_{x_k^E} P(x_{k+1}^E | x_k^E, a_{k+1}^R, al_{k+1}, a_k^O) P(x_k^E | \text{evidence}_{1:k}) \quad (6)$$

This enables forward simulation of potential future states, allowing the model to anticipate user behavior under different interaction scenarios.

4 Observational Study and Model Learning

We adopt a data-driven approach to refine the relationships in our model by conducting an observational study in a controlled setting. This empirical study allow us to analyze the dependencies between latent states and observed actions, providing insights that helped parameterize our DBN models based on real-user data.

Once the structural relationships are established, we parameterize the model by estimating the conditional probability distributions (CPDs) for each variable given its parent nodes. To achieve this, we employ Bayesian parameter estimation, incorporating prior knowledge through a Dirichlet prior with a uniform hyperparameter α .

4.1 Observational Study

In our observational study, we focus on the interaction between a self-driving scooter as the AV and delivery robots as other road users. We investigate accommodative actions at the strategic level, consisting of yielding (as positive prosocial action) and unyielding actions during a conflict of path in an interaction. We specifically chose to examine sidewalk interactions, where there are no formal road rules regarding yielding or unyielding actions. This setting allows us to explore how these road users navigate shared spaces without predefined rules, relying instead on their accommodative actions when faced with potential path conflicts.

Experiment Design. We conducted a mixed design study, measuring the impact of delivery robot and scooter accommodative actions on user's well-being and trust. The study was designed as a $2(\text{Other's action}) \times 2(\text{Ego's action})$ experiment, where the user rode a self-driving scooter and interacted with a delivery robot during the ride. Other's action had two levels: the robot yielding to the ego O_+ or unyielding to the ego O_- . Similarly, ego's action had two levels: the ego scooter yielding to the delivery robot R_+ or unyielding to the delivery robot R_- . This study was conducted using a Wizard of Oz experiment, where human operators controlled the scooter's movements to simulate strategic decision-making processes. This approach allowed us to focus on analyzing yielding and unyielding behaviors at the action level without the influence of specific trajectory-planning algorithms.

Events and Scenarios. Each participant interacts with two rides, with each ride comprising a sequence of two events: first event with delivery robot as the contributor (ego as the receiver) and second event with ego as the contributor (delivery robot as a receiver). This two-event sequence repeats in the second ride, with different scenarios. Four interaction scenarios were developed where the contributor's accommodative action is either yielding or unyielding. Figure 2 displays the four scenarios ($S1-S4$) used in the study. The yielding actions in these scenarios only include actions such as stopping for the other to go first ($S1$ and $S3$), merging ($S4$), or changing the way for the other ($S2$). The order of scenarios across the four events was counterbalanced based on a Balanced Latin Square design.

We have two independent variables with two levels each: robot's accommodative actions (O_+ and O_-) and ego's accommodative actions (R_+ and R_-). These variables are ma-




Figure 2: Four scenarios of ego-delivery robot interaction. In scenarios $S1$, $S2$, and $S3$, the roles of contributor and receiver are interchanged depending on who takes the yielding or unyielding action first. In scenario $S4$, the scooter acts as the contributor while the robot is the receiver. The robot and scooter switch locations in the robot-contributor scenario of $S4$.



Figure 3: Web-based riding environment.

nipulated as both within-subjects and between-subjects factors. Given that each ride consists of an accommodative action by the robot followed by accommodative action by the scooter, there are 4 (2×2) possible accommodative action combinations. To reduce the number of cases, the second ride has the same accommodative action from the robot as in the first ride; the scooter has all four permutations of accommodative actions across the two rides. This results in a total of 32 combinations. After each event, we measure the user’s well-being and trust using a self-report questionnaire. During the second event in each ride, where the ego acts as a contributor, the user’s intention toward the delivery robot is asked before the scooter exhibits its accommodative action. This allows to determine whether the user’s intentions align with the scooter’s action or not.

Well-being and Trust Questionnaire. To assess user’s well-being, we used a modified version of the social well-being questionnaire [Radzyk, 2014] and made it situational rather than general and more applicable to our study scenario. The questionnaire was designed to measure well-being based on four factors: (1) positive relationship, (2) satisfaction with travel, (3) trust, and (4) general well-being. To calculate user’s well-being, we average across all seven questions. Additionally, we asked a question specifically related to the user’s trust in the self-driving scooter to measure trust independently (The details of the questionnaire used in this study are in the supplementary material.) To ensure a more consistent users’ understanding of well-being and trust, we defined these concepts during the introduction of the user study to the participants. Moreover, users had access to the definitions of specific words (such as ‘trust’) while answering the questions.

User’s intention. To determine user’s intention in the interaction where the ego is the contributor, we ask them a specific question: “What action would you like your self-driving scooter to take regarding the delivery robot?” Two options

are given, with one implying yielding action and the other unyielding action. Based on the user’s response, we assess whether the user’s intention aligns with the action of the self-driving scooter in that particular event.

Participants. A total of 300 participants were recruited via Prolific (<https://www.prolific.co/>), with the majority of participants being between 25 and 55 years old (71.23%). Of the participants, 54% identified as males, 44% as females, and 2% as others. All participants passed attention checks and indicated their commitment to thoughtfully answering survey questions. Participants were compensated \$3.0 for their approximately 25-minute participation. All participants provided informed consent. The study was approved by the Bioethics Committee in Honda R&D (approval code: 99HM-065H)

Stimuli. The study was conducted online, and we used video recordings from a custom medium-fidelity driving simulator to simulate the scenarios. The simulated environments were created using Unreal Engine 4.27 (<https://www.unrealengine.com/>) with AirSim [Shah *et al.*, 2018]. To provide a realistic experience, the videos were recorded using one front-facing camera with 133 degree horizontal field of view. Additionally, the scooter’s speed information was overlaid at the bottom of the screen to enhance the experience (see Figure 3). The details of the study procedure, along with example recordings of events and scenarios are provided in the supplementary material. The supplementary materials, including the dataset and source code for the models discussed in the following sections, are available on GitHub at <https://github.com/honda-research-institute/wellbeing-trust-model>.

4.2 Refinement of State Relationships

Using the collected data, we examined statistical dependencies between latent states and observed variables to refine the structural relationships within the DBN models. Specifically, we applied statistical one-tail and two-tail t-tests to assess conditional independencies and validate the initial model structure. Across all 32 conditions of the study, we grouped the participants’ data into different groups based on three factors: (1) robot’s yielding vs. unyielding action, (2) scooter’s yielding vs. unyielding action, and (3) alignment of participant’s intention with the scooter’s action.

Data analysis was conducted to identify correlations between the scooter’s action a_k^R , the user’s intention alignment state al_k , robot delivery’s action a_k^O , and the evolution of the user’s

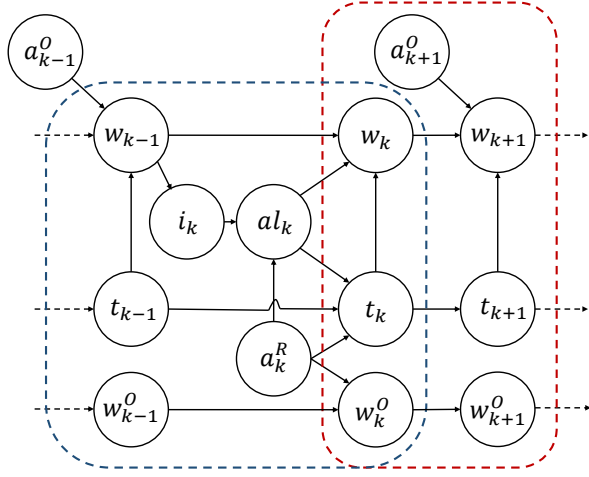


Figure 4: Final model structure of R-DBN (blue) and O-DBN (red)

latent states x_k^E . Among these, we identified statistically significant correlations with a p-value of < 0.05 . Specifically, we have the user’s well-being is higher ($t(594.87) = 7.65$, $p < 0.0001$ ****) when the delivery robot exhibits yielding action toward the ego compared to unyielding action. The user’s trust is higher when the scooter exhibits yielding action than unyielding action toward others ($t(576.114) = 5.54$, $p < 0.0001$ ****). The users with yielding intention toward others has higher well-being than those with unyielding intention ($t(317.49) = 2.12$, $p = 0.02$). The user’s well-being and their trust in the scooter are higher when the scooter’s action toward others is aligned with the user’s intention compared to when it is not aligned ($t(577.26) = 7.64$, $p < 0.0001$ ****). User’s trust is positively correlated with well-being ($r(597) = 0.7058$, $p < 0.0001$ ****). Based on these findings, we adjusted the structure of the DBN models by setting the likelihood of relationships between variables without significant correlation to zero. To capture other’s well-being, we assume that the effect of scooter’s action on robot is the same as robot’s actions on the user. To ensure that the model accurately reflects the observed dependencies in the data we adjusted the causal structure of the model. In particular, we performed a 5-fold cross-validation to evaluate candidate models and selected the best structure based on log-likelihood scores (SC-DBN: -2854.90 and RC-DBN: -2698.20). See Figure 4 for final model structure.

4.3 Bayesian Parameter Estimation

After finalizing the structure of the DBN models, we estimated the conditional probability distributions (CPDs) for each node using Bayesian parameter estimation. Given the dataset of observed state transitions, we computed the posterior distribution over CPDs using a Dirichlet prior.

To facilitate inference, we queried the propagated belief at event k based on past experiences. To simplify CPDs, we scaled latent state values between 0 and 1 and discretized them into six bins (selected through ablation studies), allow-

ing representation in a tabular format.

For discrete variables, the CPDs were represented as tabular distributions, where the conditional probability of a variable X given its parent variables $Pa(X)$ is given by

$$P(X|Pa(X)) = \frac{N + \alpha}{\sum_X N + \alpha.k} \quad (7)$$

where N is the observed count of transitions from the dataset, α is the Dirichlet prior, and k is the number of possible states for each variable.

Using our dynamic Bayesian networks and available data, we estimated these CPDs with the Bayesian Parameter Estimator, applying a Dirichlet prior with a uniform α . We implemented our DBN models using the pgmpy library in Python [Ankan and Panda, 2015].

By leveraging Bayesian parameter estimation and discretization, our DBN effectively captures probabilistic dependencies while enabling robust inference across different interaction scenarios.

4.4 Model Evaluation

To evaluate the model’s performance in inferring key variables, we focused on assessing its accuracy in predicting user’s *well-being*, *trust*, and *intention*. Accuracy was measured as the proportion of correct inferences out of all the inferences made during the evaluation. We employed 5-fold cross-validation to assess the model’s generalization performance, repeating the cross-validation procedure for 100 iterations to ensure stable and reliable accuracy estimates.

The model achieved an accuracy of 77% for inferring *well-being*, 67% for *trust*, and 95% for *intention*, which demonstrates its effectiveness in capturing the dynamics of these important variables.

In addition to accuracy, we also evaluated how well the model could infer changes in these variables over time under different conditions. Specifically, we examined the evolution of *well-being* w_k , *trust* t_k , and *other’s well-being* w_k^O during 10 consecutive events of scooter interactions, where we fixed evidence and control variables such as *action alignment* and the *scooter’s action*. This approach allowed us to isolate the effects of action alignment and the scooter’s actions on the inferred variables. To further assess the model’s performance, we analyzed how the inferred variables evolved across these 10 events. Figure 5 illustrates the expected values for the user’s *well-being* w_k , *trust* t_k , and *other’s well-being* w_k^O over these events. As seen in the figure, the user’s well-being increases over time when the scooter’s action is aligned with the user’s intention and decreases otherwise. Trust also shows dynamic changes depending on the action alignment and accommodative actions. Furthermore, other’s well-being increases when the scooter takes yielding action and decreases when it takes unyielding action.

These findings highlight the model’s capacity to infer and capture the dynamic relationships between the variables and the actions taken by the scooter, providing valuable insights into user behavior over time.




Figure 5: Expected values of inferred states over 10 events using the model given the scooter's actions and action alignment.

5 Informed Decision Making

One of the primary applications of the proposed model is to develop human-aware automation that can account for human cognitive states while making AV decisions. In this section, we will focus on the use of our model to determine the optimal policy or action for a self-driving scooter (AV) despite the uncertainty in the user's state variables.

For optimal decisions, we need to define a utility function that we want to maximize given the model. We use causal inference modeling (CIM) to determine a policy that the self-driving scooter can adopt to optimize certain factors while taking uncertainty into account. CIM is a generalization of the Bayesian network that is used to represent decision-making processes under uncertainty. We transform our proposed dynamic Bayesian models into CIM with a casual inference diagram with chance nodes, utility nodes, and decision nodes. A utility node represents the outcome or value that a decision-maker is trying to optimize. A chance node represents an uncertain event or variable that can affect the outcome, while a decision node represents a point in the diagram where a decision-maker has control over the value of a variable. Depending on the factor that the scooter aims to optimize, the utility node may differ, such as user's well-being, user's trust, other's well-being, scooter's costs, or a trade-off between multiple factors. Additionally, when the generated policy pertains to the actions the scooter should take, the decision node represents the scooter's accommodative action, denoted as a_k^R . We used PyCID library in Python [James Fox *et al.*, 2021] to implement our casual inference modeling. We used the conditional probability distributions of each variable that we estimated using the data for our dynamic Bayesian models to build the CIM.

Given our proposed model and the utility function we want to optimize, we can use casual inference modeling to reason over the expected utility of each action and choose the action that maximizes the expected utility. Formally, given the evidence ev , the policy π at event k is given by

$$\pi = \arg \max_{a_k^R \in \{E_u, E_y\}} \{E[U(a_k^R|ev)]\}. \quad (8)$$

Here, $E[U(a_k^R|ev)]$ represents the expected utility of taking the action a_k^R , which is computed as the sum over probability of all possible outcome states of a_k^R , i.e., $O_i(a_k^R)$, given the evidence and the action, and then multiplied by its corresponding utility function $U(O_i(a_k^R)|a_k^R)$ as

$$E[U(a_k^R|ev)] = \sum_i p(O_i(a_k^R)|ev, a_k^R) U(O_i(a_k^R)|a_k^R) \quad (9)$$

By defining an appropriate utility function, we can find the action that maximizes the expected utility. We analyze differ-

ent policies that the scooter can adopt based on its objectives: (1) maximizing user's well-being ($U_k = w_k$), (2) maximizing user's trust ($U_k = t_k$), and (3) optimizing a trade-off of user's well-being, other's well-being, and the cost of its actions ($U_k = w_k + w_k^O + C(a_k^R)$). With these utilities, we can determine the optimal actions the scooter should take in different scenarios to achieve its desired objective.

Maximizing User's Well-being. We define the utility function as equal to the value of well-being. Using the model, we then determine the policy that maximizes user's well-being. In the absence of any evidence, the optimal policy for the scooter is to always take a yielding action. To further analyze the policy, we conducted an analysis of which nodes have a positive value of information (VOI) incentives. VOI incentives represent the expected increase in utility that can be achieved by acquiring additional information. VOI analysis shows that having evidence of previous well-being (w_{k-1}), trust (t_{k-1}), and current user intention (i_k) can lead to better reasoning of the optimal policy and consequently improve the utility. Thus, we analyze the optimal policy given the availability of evidence on these variables. Results show that if the scooter has access to information about the previous user's well-being, the optimal policy would be

$$a_k^R = \begin{cases} R_- & \text{unyielding} & \text{if } 0 \leq w_{k-1} \leq 0.18 \\ R_+ & \text{yielding} & \text{otherwise} \end{cases}. \quad (10)$$

The optimal policy with information about the user's intention is to be aligned with user's intention, i.e.,

$$a_k^R = \begin{cases} R_- & \text{unyielding} & \text{if } i_k = I_- \\ R_+ & \text{yielding} & \text{if } i_k = I_+ \end{cases}. \quad (11)$$

Furthermore, the availability of previous trust information does not alter the policy when no evidence is available.

Maximizing User's Trust. We define the utility at event k as the user's trust at that event. We identify that the previous user's well-being w_{k-1} , trust t_{k-1} , and current user's intention i_k have positive VOI when the scooter aims to maximize user trust. When no evidence is available, the optimal policy is always to take the yielding action. However, if information about the user's intention is available, adopting the alignment policy can increase trust (similar to the case $U_k = w_k$), while having information about user's previous well-being and trust does not affect the optimal policy when no evidence is available.

Optimizing the Trade-off between User's Well-being, Other's Well-being and Scooter's Cost. In this scenario, we consider that the cost of the scooter's action is different




Figure 6: Optimal accommodative policy at event k . Sensitivity study results based on costs and information variables w_{k-1} .

depending on whether it takes a yielding or an unyielding action. Since yielding actions usually involve changing direction or stopping, they are more costly (for example, more fuel spent) than unyielding actions. To simplify the problem, we assume that the cost difference is incorporated in the utility function, which becomes:

$$U_k = \begin{cases} w_k + w_k^O & \text{if } a_k^R = R_- \\ w_k + w_k^O + C & \text{if } a_k^R = R_+ \end{cases} \quad (12)$$

The analysis shows that the cost of actions can significantly impact the utility and, thus, the policy. Therefore, we perform a sensitivity analysis of the cost values and examine how the policy changes according to different costs. We consider costs to be any value greater than or equal to zero. However, since the values of well-being range between 0 and 1, we find that the most critical costs are the ones that fall within this range. Figure 6 illustrates how the policy changes with different information on the user’s well-being. Analysis over other positive VOI evidence, including, trust, intention, and others’ well-being, as well as in the absence of evidence shows that for costs less than 0.2, the scooter should take yielding action and for more than that should take unyielding action.

To sum up, we can see that if the cost difference between the two types of actions is large, the optimal action for the scooter is to always take unyielding actions. However, this would sacrifice the well-being of the user and others. Thus, when the cost difference is lower, we must balance the trade-off between well-being and cost to determine the optimal policy. For a cost difference of approximately 0.2, the scooter’s optimal policy is to mostly take yielding actions. However, for costs between 0.2 and 0.4, the scooter should mostly take unyielding actions unless the user or others are in a very high well-being state. Therefore, our proposed causal inference model allows us to make informed decisions about the best course of action for a scooter in order to achieve its objectives while balancing various factors.

6 Conclusions and Future Work

Our paper proposes a computational model for autonomous vehicles that infers the cognitive states such as well-being, trust, and intentions of users and other road user’s well-being

to make informed decisions. Our DBN models provide a structured representation of these states, enabling probabilistic inference. To learn the parameters of this model, we conducted an observational study to collect interaction data, which was then used to refine state relationships, estimate Bayesian parameters, and evaluate the model’s effectiveness. By integrating this model into AV decision-making, the system can account for well-being, trust, and cost trade-offs, leading to safer and more user-centered interactions with potential to inspire future research in autonomous vehicles.

Finally, while our proposed model provides a structured framework for inferring trust, well-being, and intention in human-AV interactions, the observational study conducted for parameter learning introduces certain limitations. Since the study consisted of only four interaction events per participant, the collected data may not fully capture the long-term evolution of trust and well-being. Future work will explore longitudinal data collection to better model these temporal dynamics. Additionally, our study focused on dyadic interactions between a self-driving scooter (AV) and delivery robots (other road users) in sidewalk environments where no formal yielding rules exist. While this setting allowed for controlled observation of accommodative actions, future research will extend the model’s applicability by incorporating interactions involving multiple road users and different mobility modes. Despite these study-specific constraints, our work demonstrates a step toward human-aware automation. By leveraging probabilistic inference, the proposed model enables AVs to reason about trust and well-being, ultimately contributing to safer and more user-centered autonomous systems. Future work can explore adapting the model by incorporating other domain-specific variables and contextual information. To enhance versatility, further research can investigate adjusting input variables and data collection methods to suit different interaction domains, accommodating a broader range of decision-making processes and cognitive states.

Acknowledgements

We would like to acknowledge Hiu-Chun Lo at Honda Research Institute USA, Inc for his support for the simulator and website development.

References

- [Akash *et al.*, 2020] Kumar Akash, Neera Jain, and Teruhisa Misu. Toward adaptive trust calibration for level 2 driving automation. In *ICMI*, pages 538–547, 2020.
- [Ankan and Panda, 2015] Ankur Ankan and Abinash Panda. pgmpy: Probabilistic graphical models using python. In *SCIPY*. Citeseer, 2015.
- [Azevedo-Sa *et al.*, 2020] Hebert Azevedo-Sa, Suresh Kumar Jayaraman, X Jessie Yang, Lionel P Robert, and Dawn M Tilbury. Context-adaptive management of drivers’ trust in automated vehicles. *RA-L*, 5(4):6908–6915, 2020.
- [Bowen and Smith, 2019] Louise Bowen and Andrew Smith. Drive better, feel better: Examining associations between

- well-being and driving behaviour in students. *Advances in Social Sciences Research Journal*, 6(2), 2019.
- [Deo and Trivedi, 2019] Nachiket Deo and Mohan M Trivedi. Looking at the driver/rider in autonomous vehicles to predict take-over readiness. *IEEE T-IV*, 5(1):41–52, 2019.
- [Ettema *et al.*, 2011] Dick Ettema, Tommy Gärling, Lars Eriksson, Margareta Friman, Lars E Olsson, and Satoshi Fujii. Satisfaction with travel and subjective well-being: Development and test of a measurement tool. *Trans. Research Part F: Traffic Psychology and Behaviour*, 14(3):167–175, 2011.
- [Friman *et al.*, 2013] Margareta Friman, Satoshi Fujii, Dick Ettema, Tommy Gärling, and Lars E Olsson. Psychometric analysis of the satisfaction with travel scale. *Transportation Research Part A: Policy and Practice*, 48:132–145, 2013.
- [Guhe *et al.*, 2005] Markus Guhe, Wayne D Gray, Michael J Schoelles, Wenhui Liao, Zhiwei Zhu, and Qiang Ji. Non-intrusive measurement of workload in real-time. In *HFES*, volume 49, pages 1157–1161. SAGE Publications Sage CA: Los Angeles, CA, 2005.
- [Guo *et al.*, 2020] Yaohui Guo, Chongjie Zhang, and X Jessie Yang. Modeling trust dynamics in human-robot teaming: A bayesian inference approach. In *CHI-Extended Abstracts*, pages 1–7, 2020.
- [Halkola *et al.*, 2019] Eija Halkola, Lauri Lovén, Marta Cortes, Ekaterina Gilman, and Susanna Pirttikangas. Towards measuring well-being in smart environments. In *UbiComp/ISWC*, pages 1166–1169, 2019.
- [Harris *et al.*, 2014] Paul B Harris, John M Houston, Jose A Vazquez, Janan A Smither, Amanda Harms, Jeffrey A Dahlke, and Daniel A Sachau. The prosocial and aggressive driving inventory (padi): A self-report measure of safe and unsafe driving behaviors. *Accident Analysis & Prevention*, 72:1–8, 2014.
- [Hu *et al.*, 2013] Tian-Yi Hu, Xiaofei Xie, and Jie Li. Negative or positive? the effect of emotion and mood on risky driving. *TRANSPORT RES F-TRAF*, 16:29–40, 2013.
- [Isler and Newland, 2017] Robert B Isler and Samantha A Newland. Life satisfaction, well-being and safe driving behaviour in undergraduate psychology students. *TRANSPORT RES F-TRAF*, 47:143–154, 2017.
- [James Fox *et al.*, 2021] James Fox, Tom Everitt, Ryan Carey, Eric Langlois, Alessandro Abate, and Michael Wooldridge. PyCID: A Python Library for Causal Influence Diagrams. In Meghann Agarwal, Chris Calloway, Dillon Niederhut, and David Shupe, editors, *SCIPY*, pages 43 – 51, 2021.
- [Koppol *et al.*, 2021] Pallavi Koppol, Henny Admoni, and Reid G Simmons. Interaction considerations in learning from humans. In *IJCAI*, pages 283–291, 2021.
- [Lee and See, 2004] John D. Lee and Katrina A. See. Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1):50–80, 2004.
- [Liang and Lee, 2014] Yulan Liang and John D Lee. A hybrid bayesian network approach to detect driver cognitive distraction. *Transp. Res. Part C Emerg. Technol.*, 38:146–155, 2014.
- [Liang *et al.*, 2007] Yulan Liang, John D Lee, and Michelle L Reyes. Nonintrusive detection of driver cognitive distraction in real time using bayesian networks. *Transportation research record*, 2018(1):1–8, 2007.
- [Luo *et al.*, 2019] Ruikun Luo, Yifan Wang, Yifan Weng, Victor Paul, Mark J Brudnak, Paramsothy Jayakumar, Matt Reed, Jeffrey L Stein, Tulga Ersal, and X Jessie Yang. Toward real-time assessment of workload: A bayesian inference approach. In *HFES*, volume 63, pages 196–200. SAGE Publications Sage CA: Los Angeles, CA, 2019.
- [Luo *et al.*, 2021] Ruikun Luo, Yifan Weng, Yifan Wang, Paramsothy Jayakumar, Mark J Brudnak, Victor Paul, Vishnu R Desaraju, Jeffrey L Stein, Tulga Ersal, and X Jessie Yang. A workload adaptive haptic shared control scheme for semi-autonomous driving. *Accident Analysis & Prevention*, 152:105968, 2021.
- [Mahmood *et al.*, 2024] Syed Hasan Amin Mahmood, Zhuoran Lu, and Ming Yin. Designing behavior-aware ai to improve the human-ai team performance in ai-assisted decision making. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 3106–3114, 2024.
- [Mehrotra *et al.*, 2023] Shashank Mehrotra, Zahra Zahedi, Teruhisa Misu, and Kumar Akash. Wellbeing in Future Mobility: Toward AV Policy Design to Increase Wellbeing through Interactions. In *ITSC*, 2023.
- [Nouvian, 2023] Tom Nouvian. In paris referendum, 89% of voters back a ban on electric scooters, Apr 2023.
- [Ong *et al.*, 2019] Desmond C Ong, Jamil Zaki, and Noah D Goodman. Computational models of emotion inference in theory of mind: A review and roadmap. *TOPICS*, 11(2):338–357, 2019.
- [Radzyk, 2014] JJ Radzyk. Validation of a new social well-being questionnaire. B.S. thesis, University of Twente, 2014.
- [Sauer *et al.*, 2019] Vanessa Sauer, Alexander Mertens, Verena Nitsch, and Jens Dietmar Reuschel. An empirical investigation of measures for well-being in highly automated vehicles. In *AutomotiveUI: Adjunct Proceedings*, pages 369–374, 2019.
- [Seligman and others, 2002] Martin EP Seligman et al. Positive psychology, positive prevention, and positive therapy. *Handbook of positive psychology*, 2(2002):3–12, 2002.
- [Shah *et al.*, 2018] Shital Shah, Debadeepta Dey, Chris Lovett, and Ashish Kapoor. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics: Results of the 11th International Conference*, pages 621–635. Springer, 2018.
- [Soh *et al.*, 2020] Harold Soh, Yaqi Xie, Min Chen, and David Hsu. Multi-task trust transfer for human-robot interaction. *IJRR*, 39(2-3):233–249, 2020.

- [Stocker and Shaheen, 2017] Adam Stocker and Susan Shaheen. Shared automated vehicles: Review of business models. International Transport Forum Discussion Paper 2017-09, Paris, 2017.
- [Wu *et al.*, 2022] Tong Wu, Enna Sachdeva, Kumar Akash, Xingwei Wu, Teruhisa Misu, and Jorge Ortiz. Toward an adaptive situational awareness support system for urban driving. In *IV*, pages 1073–1080. IEEE, 2022.
- [Xu and Dudek, 2015] Anqi Xu and Gregory Dudek. Optimo: Online probabilistic trust inference model for asymmetric human-robot collaborations. In *HRI*, pages 221–228. IEEE, 2015.
- [Xu and Dudek, 2016] Anqi Xu and Gregory Dudek. Maintaining efficient collaboration with trust-seeking robots. In *IROS*, pages 3312–3319. IEEE, 2016.
- [Zahedi *et al.*, 2022] Zahra Zahedi, Sarath Sreedharan, Mudit Verma, and Subbarao Kambhampati. Modeling the interplay between human trust and monitoring. In *HRI*, pages 1119–1123. IEEE, 2022.
- [Zahedi *et al.*, 2023] Zahra Zahedi, Mudit Verma, Sarath Sreedharan, and Subbarao Kambhampati. Trust-aware planning: Modeling trust evolution in iterated human-robot interaction. In *HRI*, pages 281–289, 2023.
- [Zhuge and Zhuge, 2020] Hai Zhuge and Hai Zhuge. Toward cyber-physical-social science. *Cyber-Physical-Social Intelligence: On Human-Machine-Nature Symbiosis*, pages 341–351, 2020.