

Improving Trust Estimation in Human-Robot Collaboration Using Beta Reputation at Fine-Grained Timescales

Resul Dagdanov, Milan Andrejević, Dikai Liu, and Chin-Teng Lin

Abstract—When interacting with each other, humans adjust their behavior based on perceived trust. To achieve similar adaptability, robots must accurately estimate human trust at sufficiently granular timescales while collaborating with humans. Beta reputation is a popular way to formalize a mathematical estimation of human trust. However, it relies on binary performance, which updates trust estimations only after each task concludes. Additionally, manually crafting a reward function is the usual method of building a performance indicator, which is labor-intensive and time-consuming. These limitations prevent efficient capture of continuous trust changes at more granular timescales throughout the collaboration task. Therefore, this letter presents a new framework for the estimation of human trust using beta reputation at fine-grained timescales. To achieve granularity in beta reputation, we utilize continuous reward values to update trust estimates at each timestep of a task. We construct a continuous reward function using maximum entropy optimization to eliminate the need for the laborious specification of a performance indicator. The proposed framework improves trust estimations by increasing accuracy, eliminating the need to manually craft a reward function, and advancing toward the development of more intelligent robots.¹

Index Terms—Beta reputation system, human trust, human-robot collaboration, probabilistic model.

I. INTRODUCTION

HUMAN decisions are often influenced by their perceptions of how trustworthy they are perceived by others [1], [2]. Research in human-robot collaboration (HRC) indicates that when robots act in accordance with a human co-worker’s trust, collaboration effectiveness is enhanced [3]–[5]. However, to make trust-aware decisions, a robot needs to accurately estimate how much its co-worker trusts it [6], [7].

Trust in a robot can change throughout a task, making it essential for the robot to estimate trust in real-time at fine-grained timescales. By continuously estimating trust during

This work was supported in part by Australian Research Council (ARC) Research Hub under Grant IH240100016 and in part by the University of Technology Sydney. (Corresponding author: Resul Dagdanov.)

Resul Dagdanov and Dikai Liu are with the Robotics Institute, Faculty of Engineering and Information Technology, University of Technology Sydney, Ultimo, NSW 2007, Australia. Resul.Dagdanov@uts.edu.au, Dikai.Liu@uts.edu.au

Milan Andrejević is with the Psychology Discipline, Graduate School of Health, Faculty of Health, University of Technology Sydney, Ultimo, NSW 2007, Australia. Milan.Andrejevic@uts.edu.au

Chin-Teng Lin is with the Australian Artificial Intelligence Institute, School of Computer Science, Faculty of Engineering and Information Technology, University of Technology Sydney, Ultimo, NSW 2007, Australia. Chin-Teng.Lin@uts.edu.au

¹The source code and media materials are publicly available at <https://github.com/resuldagdanov/robot-learning-human-trust>

©2026 IEEE

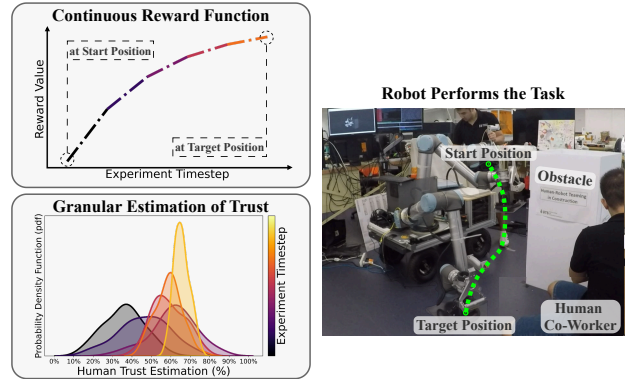


Fig. 1. Overview of a real-time estimation of human co-worker’s trust. The objective of the intelligent robot is to autonomously transport tiles from a randomly chosen starting position to a target position. While performing the task, the framework continuously assigns a reward value at each timestep. The proposed framework mathematically links the reward values to the probabilistic trust estimation at fine-grained timescales. Doing so enables a real-time estimation of human trust at each timestep throughout the task.

the task rather than only at its conclusion, the robot can adapt its behavior immediately, either enhancing or reducing trust to address the pitfalls of overtrust or undertrust [8]–[10].

There is growing HRC research interest in computational models to estimate human trust toward robots [11]–[14]. These models are based on robot performance, which is the most significant factor influencing human trust [10], [15], [16]. Furthermore, probabilistic models that capture uncertainty and bias in human subjectivity show great promise in this context [17]–[20]. The framework in this letter entails a probabilistic estimation of trust based on performance (Fig. 1).

The probabilistic models proposed in [17]–[20] fail to capture the continuous changes in human trust as a robot performs a task. This limitation arises because human co-workers assess performance in a binary manner (e.g., success or failure) only after task completion, neglecting performance changes during the task. This results in a static estimation of trust dynamics, often referred to as a “snapshot” view [19].

Formulating a task-specific performance function is labor-intensive, time-consuming, and requires a deep understanding of the task [21]–[23]. This process involves determining appropriate weights for the factors influencing task objectives and aligning them with desired outcomes. This limits the autonomy and adaptability of the robot to various tasks and highlights the need for a framework that enables fine-grained estimation of human trust, facilitating real-time trust-aware robot decision-making with minimal labor-intensive effort.

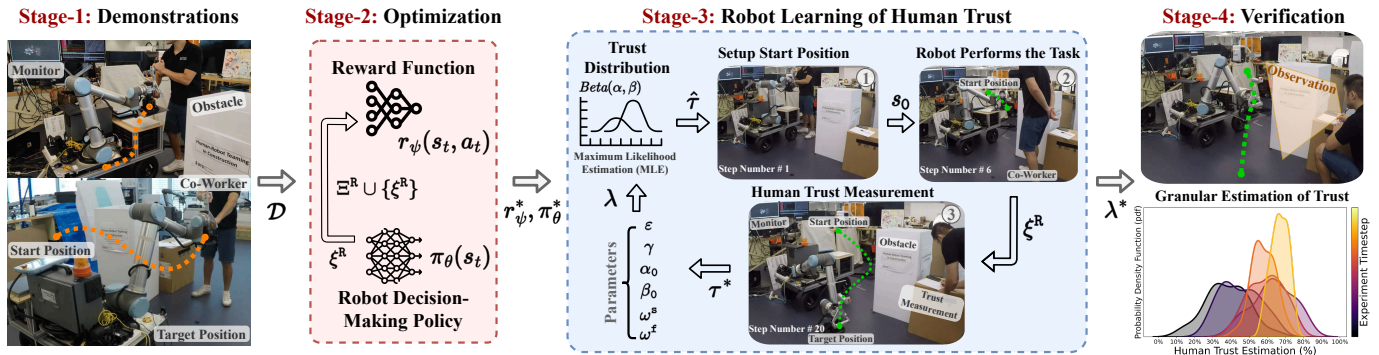


Fig. 2. **Proposed framework for human trust estimation in HRC.** **Stage-1:** A co-worker demonstrates the task objectives to an intelligent robot by physically moving the robotic arm. **Stage-2:** In this stage, a reward function, r_{ψ} , and the robot’s decision-making policy, π_{θ} , are optimized through maximum entropy optimization and behavior cloning, respectively. **Stage-3:** The robot iteratively learns the trust dynamics of the human co-worker over multiple consecutive tasks. For each task, the co-worker randomly sets a starting position for the robotic arm. The robot performs the task using the optimized decision-making policy (π_{θ}). After the robot completes the task ($T = 20$), the co-worker self-reports trust. The optimization objective is to minimize the difference between the measured trust (τ) and the estimated trust ($\hat{\tau}$) using maximum likelihood estimation. The parameters for optimization include success and failure weights, initial probability distribution parameters, the success-failure reward determination threshold, and the trust history dependency constant, collectively represented as $\lambda = \omega^s, \omega^f, \alpha_0, \beta_0, \epsilon, \gamma$. **Stage-4:** In the verification stage, the robot estimates the human co-worker’s trust in real-time at fine-grained timescales.

We propose a new framework for the accurate estimation of human trust at granular timescales, as visualized in Fig. 2. We construct a continuous reward function using maximum entropy optimization, which enables us to efficiently capture the underlying performance dynamics throughout the task.

This study focuses on trust estimation at fine-grained timescales. The proposed framework establishes a mathematical connection between continuous reward functions and probabilistic trust dynamics. Although the current framework does not yet incorporate robot behavior adaptation based on estimated trust, we consider fine-grained trust estimation a critical foundation for enabling real-time behavioral adaptation. To elaborate, in Section VI, we outline potential strategies for trust-aware decision-making. For instance, if a human’s trust declines following a task failure, the robot could communicate an apology and a commitment to improve [9], [24]. Such communications may help mitigate undertrust in real-time, preventing disproportionate deterioration. Developing these behavioral adaptations requires further research, which we leave for future exploration.

II. LITERATURE REVIEW

A. Probabilistic and Deterministic Trust Models in HRC

A computational model is necessary for the estimation of human trust. One straightforward way is to construct trust estimation as a linear combination of performance features that influence trust in HRC [13], [25]. Such models do not incorporate uncertainty, which is advantageous in some tasks. For example, in repetitive assembly tasks where the robot’s reliability is consistent and predictable, deterministic models of trust are applied [25], [26]. However, in cases where it is important to capture human subjectivity, these models fall short. This is because human perceptions and decision-making typically involve uncertainty [1], [2], [27]. Therefore, human perceptions of trustworthiness most likely include subjective uncertainty, which informs their decisions.

Similarly, for robots to adjust their behavior based on human trust, they must capture uncertainty in their estimations.

To capture uncertainty in trust estimations, it is required to apply probabilistic models. A typical model is a dynamic Bayesian network. However, this model lacks a mathematical framework to describe how human trust stabilizes over time through repeated collaborations with the same robot [28]. A more suitable alternative is the beta reputation, which has been proposed to address this limitation [19], where a beta distribution offers two main advantages [29]. Firstly, it limits an estimation interval to 0 and 1, creating consistency with the trust measurement scale. Secondly, this model accounts for a historical reputation by accumulating the number of successful and unsuccessful collaborations. These advantages make beta reputation suitable for the probabilistic estimation of trust.

B. Robot Learning of Human Trust From Demonstrations

To reduce the laborious workload in designing trust models, effective use of demonstrations is required. These demonstrations provide insight into how a co-worker expects a robot to perform tasks [30]. Note that trust dynamics in HRC closely depend on the co-worker’s expectations of the robot’s capabilities [15], [31]. Thus, these demonstrations are a critical resource for modeling trust, as they reflect human expectations.

The maximum entropy optimization method can quantify the similarity mismatch between co-worker demonstrations and a robot’s capabilities [32]. It was applied to construct reward and trust estimation models using demonstrations [18]. However, by clustering similar states into a fixed number of decision-making policies, this method loses flexibility in environments where decision-making parameters are not constant. Furthermore, recent work in [20] has applied this method to learn personalized weights in trust estimation. Similarly, the work in [33] has used the preference-based optimization [34] to capture individual differences. For example, one co-worker may prioritize safety, while another may emphasize a robot’s speed as a critical trust indicator.

III. PROBLEM FORMULATION

A. Markov Decision Process

Let us consider a decision-making policy where an intelligent robot performs the task in alignment with human demonstrations. The expression for this process is a Markov decision process (MDP), denoted as $\mathcal{M} := \langle \mathcal{S}, \mathcal{A}, r, f, \mathbf{T} \rangle$. At each timestep (step number) t in a task that concludes after a total of T timesteps, $s_t \in \mathcal{S}$ represents a robot state vector, and $a_t \in \mathcal{A}$ represents a robot action vector. An intelligent robot performs an action under the decision-making policy and then transitions to a new state $s_{t+1} = f(s_t, a_t)$ after receiving a reward $r(s_t, a_t)$, where f is a transition function.

B. Robot Operations and Human Demonstrations

Robot operations can take various forms, such as kinesthetic navigation, audio, and visual communication. In this work, operations performed by a robot are robotic arm manipulations under the control of a decision-making policy. As a robot performs actions, it transitions to a new state based on transition function dynamics, similar to those described in [21], [22]. A consecutive sequence of these transitions is a spatial trajectory, which is represented as a finite set of T timestep state-action pairs and notated as $\xi = \{(s_1, a_1), (s_2, a_2), \dots, (s_T, a_T)\} \in \Xi$. To simplify the notation, as adopted from [21], a trajectory can be denoted in a compact form as $\zeta = (s_1, a_1, a_2, \dots, a_T)$.

Human demonstrations refer to samples where a co-worker physically shows the robot how to perform a task by manipulating the robotic arm, like the kinesthetic interactions described in [13]. The notation for a dataset of N demonstrations is $\mathcal{D} = \{\xi_1^H, \xi_2^H, \dots, \xi_N^H\}$, where ξ^H is a sample trajectory demonstration by a human co-worker.

C. Human Trust Definition

In this letter, a human is the trustor, and a robot is the trustee. The concept of human trust is defined in [31] as “the attitude that an agent will help achieve an individual’s goals in a situation characterized by uncertainty and vulnerability”. This definition of trust aligns with the engineering aspects and the goal-oriented nature of a robot in HRC.

In HRC literature [7], models of human trust estimations are categorized into relation-based and performance-based models. Performance-based models estimate trust primarily based on the capability and reliability of a robot. In contrast, relation-based models use data on the societal and ethical norms of a human as a trust estimation feature. This letter presents a performance-based model of human trust because the engineering objective is to enhance the ability of a robot to perform physical tasks with high performance.

The performance of the robot in HRC is the most dominant factor affecting human trust in the robot [10], [15], [16]. Trust depends on the successful and unsuccessful reputation of collaboration with the robot. In this letter, a success metric of collaboration is the compatibility between human co-worker expectations and robot capabilities, which serves as a key indicator of changes in trust [30]. Based on these findings, we present a performance-based model of trust.




Fig. 3. **Measurement of human trust and numerical conversion.** We utilize a 7-point Likert scale to obtain human co-worker trust measurements, as this self-reporting technique is well-established in the literature [7], [16].

The assumption in this letter is that the reward function represents the robot’s task performance. To examine the relationship between the reward function and task performance, Section V provides an experimental analysis.

At the end of each experiment, a human co-worker self-reports their trust using a 7-point Likert scale. As shown in Fig. 3, we numerically interpret the trust values as percentages. This conversion to the percentage scale is kept constant and is symmetrical around the midpoint (i.e., 50%, which corresponds to the “Neutral” level of trust on the scale). This allows the robot to estimate trust based on numerical data.

D. Robot Decision-Making Policy

To enhance the robot’s generalization capabilities, particularly in high-dimensional action spaces, a learning-based method for robot decision-making needs to be implemented. One common approach to doing this is to maximize the similarity between expert demonstrations and robot operations, also known as behavior cloning (BC). This approach comes with a major limitation. Mimicking human actions may result in a suboptimal decision-making policy for the robot, as humans do not always act optimally. However, there is also a big advantage to this approach. As discussed in Section III-C, the similarity mismatch between co-worker demonstrations and robot operations can serve as a key indicator of changes in trust, which can be captured via the reward function. For this reason, in our framework, the human co-worker’s demonstrations are treated as optimal (i.e., expert).

A common practice in formulating a robot decision-making policy is to assume that human demonstrations resemble a Gaussian distribution [21], [23]. Consequently, a robot policy is $\pi_\theta(a_t | s_t) \sim \mathcal{N}(\mu_\theta(s_t), \sigma_\theta(s_t)^2)$, where $\mu_\theta(s_t)$ and $\sigma_\theta(s_t)$ denote the mean and variance of the policy distribution, respectively. In this manner, μ_θ and σ_θ serve as parameters of a nonlinear neural network.

$$\pi_\theta^* = \underset{\theta}{\operatorname{argmin}} \frac{1}{2} \cdot \mathbb{E}_{\substack{\xi^H \sim \mathcal{D} \\ (s_t, a_t) \sim \xi^H}} \left[\eta \cdot \log \sigma_\theta(s_t)^2 + \frac{(a_t - \mu_\theta(s_t))^2}{\sigma_\theta(s_t)^2} \right] \quad (1)$$

Eq. 1 represents a minimization objective function for optimizing a robot decision-making policy. The η coefficient acts as a scaling factor for regularizing a policy variance. When η is large, a policy optimization prioritizes minimizing a decision-making uncertainty. In contrast, when η is small, a robot performs more stochastic actions, leading to increased exploration. Section IV-A provides a detailed description of how η varies throughout this data-driven optimization process, which occurs in Stage-2 of Fig. 2.

E. Reward Function Optimization

Formulating a reward function that accurately captures the context-dependent performance of a task is laborious and time-consuming [21]–[23]. This complexity arises from the challenge of determining appropriate weights that align the reward function with desired outcomes. One solution to this problem is to use a data-driven approach, specifically maximum entropy (MaxEnt) optimization [32], to construct a reward function that accurately captures the performance of robotics applications [23], [35]. So, to eliminate the need for laborious performance specifications, we use MaxEnt optimization to construct a continuous reward function.

In the early learning epochs of BC in Stage-2 of Fig. 2, a robot performs suboptimal actions. As the optimization process in Eq. (1) proceeds, a decision-making policy gets closer to an optimal policy. The main idea behind MaxEnt optimization is to iteratively sample trajectories from $p(\xi) \sim \exp(\mathcal{R}(\xi))$. The goal is to match features between robot trajectories Ξ^R and human demonstrations \mathcal{D} .

The proposed framework does not focus on learning a robot policy by maximizing cumulative rewards. Instead, the aim of this letter is to learn a robot policy through BC by enabling it to mimic demonstrations (see Section III-D).

$$\mathcal{R}_\psi(\xi) = \frac{1}{T} \sum_{(s_t, a_t) \in \xi} r_\psi(s_t, a_t) \mid r_\psi : (\mathcal{S}, \mathcal{A}) \rightarrow \mathbb{R}^{[-1, 1]} \quad (2)$$

Formulating a linear reward function in high-dimensional environments is challenging and often impractical. A nonlinear approach is necessary to construct a reward function for these environments. A widely adopted method for capturing performance is the use of nonlinear neural networks, which provide flexibility and adaptability in data-driven solutions [23]. Therefore, we construct a reward function in Eq. (2) using a neural network with parameters ψ .

$$\mathcal{L}_{\text{MaxEnt}} = - \mathbb{E}_{\xi^H \sim \mathcal{D}} \left[\log p(\xi^H \mid \psi) \right] = - \mathbb{E}_{\xi^H \sim \mathcal{D}} \left[\log \frac{\exp(\mathcal{R}_\psi)}{\mathcal{Z}(\psi)} \right] \quad (3)$$

Eq. (3) is a loss function of MaxEnt optimization. There are infinitely many discrete possible states and actions for the background partition function $\mathcal{Z}(\psi) = \int \exp(\mathcal{R}_\psi) d\xi$ to calculate when \mathcal{S} and \mathcal{A} are both continuous.

$$\mathcal{Z}(\psi; \theta) \approx \frac{1}{M} \sum_{\xi_j^R \in \Xi^R} \left[\frac{\exp(\mathcal{R}_\psi(\xi_j^R))}{p(\xi_j^R; \theta)} \right] \quad (4)$$

A stochastic sampling-based method is a popular technique for approximating $\mathcal{Z}(\psi; \theta)$, as proposed in [23]. This method, as shown in Eq. (4), approximates close to the expectation of the negative log-likelihood loss in Eq. (3). In this letter, \mathcal{D} and Ξ^R are the sets of N demonstrations and M robot trajectories, respectively. Note that it is common practice to generate robot trajectories from $p_\theta(\xi)$ [21], [35].

To address the exploration-exploitation dilemma inherent in a decision-making policy and to enhance the generality of a reward function by approximating $\mathcal{Z}(\psi; \theta)$ in Eq. (4), this letter demonstrates an application of a dynamically annealing/interpolating linear weight η (see Eq. (1)). Section IV-A comprehensively explains the rationale behind this choice.

IV. PROBABILISTIC FRAMEWORK FOR GRANULAR ESTIMATION OF HUMAN TRUST

This section outlines the mathematical foundations for robot learning of human trust in a probabilistic manner. The divergence between human expectations and robot actions in HRC are key factors influencing the dynamics of human trust [7], [15], [30]. Accordingly, this section provides a mathematical framework for the estimation of human trust at granular timescales based on a continuous reward function.

A. Robot Decision-Making Policy and Reward Function

The denominator in Eq. (4) represents the MDP collection of $\pi_\theta(a_t \mid s_t, s_{t-1}, \dots)$, which denotes the probability of taking action a_t at state s_t according to the decision-making policy π_θ . Finding the exact value of the background partition function $\mathcal{Z}(\psi; \theta)$ with $p_\theta(\xi)$ is infeasible, especially when the ideal reward parameters ψ^* are unknown.

$$\eta = \eta_{\min} + \frac{k}{K} \cdot (\eta_{\max} - \eta_{\min}) \quad (5)$$

In [23], researchers employed an iterative method for a decision-making policy that explored the task environment by performing random actions. To minimize the frequency of the random actions, [35] investigated the use of annealing and interpolation methods to determine the exploration weight factor. The framework in this letter utilizes a linear interpolation method to find dynamic importance weight η . The method is given in Eq. (5), where k is the current epoch and K is the maximum number of learning epochs.

In order for a data-driven reward function to represent task performance, it is necessary for a robot to explore the task environment widely via random actions. During the early epochs of learning, a decision-making policy under the objective function in Eq. (5) prioritizes a wider exploration of the task environment. Consequently, η gradually increases as learning epochs increase. The selection of the hyperparameters $\eta_{\min} = 0.05$ and $\eta_{\max} = 1.00$ in Eq. (5) is based on problem-specific trials. It is important to note that increasing η_{\max} could lead to a less generalized reward function.

The objective is to gradually increase the loss function variance (uncertainty) factor during the learning process of a robot decision-making policy. This strategy ensures the robot exhibits significant uncertainty but explores more in the early policy and reward learning epochs. Adequate exploration is crucial for achieving a generalizable reward function in MaxEnt optimization. This process of learning a decision-making policy takes place in Stage-2 of Fig. 2.

The collection of human demonstration data occurs in Stage-1 of Fig. 2. Subsequently, the learning processes for the reward function and robot decision-making policy occur in Stage-2. The total number of human demonstrations remains constant throughout the proposed framework. The parameters θ^* and ψ^* represent the optimized decision-making policy and reward function obtained through MaxEnt optimization and BC, respectively. Once optimized, θ^* and ψ^* remain fixed. Therefore, for a given state s_t , the framework ensures that the reward value r_{ψ^*} is reproducible and the decision-making policy π_{θ^*} always takes action a_t .

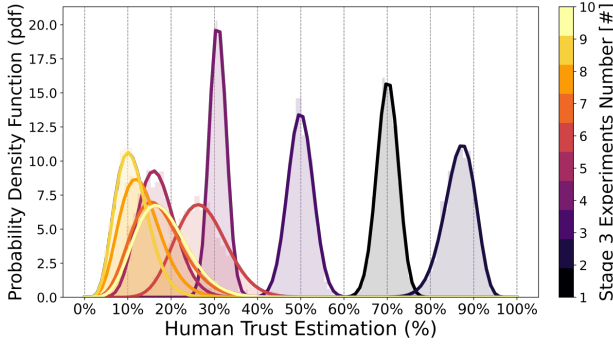


Fig. 4. **Trust estimation dynamics during Stage-3 experiments.** Trust estimations show narrower distributions in earlier experiments (darker shades), reflecting lower variance due to limited measurements. As experiments progress (lighter shades), distributions shift toward lower trust regions, and variance increases. This characteristic results from lower self-reported trust in Stage-3, which expanded the range of trust measurements (see Fig. 5).

B. Beta Reputation Model at Fine-Grained Timescales

An important theoretical part of the proposed framework is an estimation of human trust by updating a beta probability distribution at each timestep of the task.

Stage-3 of Fig. 2 involves a human co-worker testing the capability of a decision-making policy and a reward function. Each complete cycle in Stage-3 represents one experiment, after which self-reported trust is collected based on a 7-point Likert scale shown in Fig. 3. While a human co-worker reports trust at the end of each task, a reward function continuously assigns reward values to each state-action pair. This aspect enables granularity in human trust estimation.

1) *Probability Distribution:* A beta reputation system [29] is a beta probability distribution. It is a common choice to estimate trust probabilistically [17], [19], [20], as it captures subjective uncertainty and variability of human trust.

$$\tau_q(s_t, a_t) \sim \text{Beta}(\alpha_n, \beta_m) \quad (6)$$

Eq. (6) is a formulation of a continuous probability distribution to represent human co-worker's trust $\tau_q(s_t, a_t)$ at robot state s_t and robot action a_t at timestep $t \leq T$ of trajectory ξ^R . q is the total number of timesteps in all tasks.

$$\hat{\tau}_q(s_t, a_t) \doteq \mathbb{E}_{(s_t, a_t) \sim \xi^R} [\tau_q(s_t, a_t)] = \frac{\alpha_n}{\alpha_n + \beta_m} \quad (7)$$

2) *Beta Distribution Parameters:* Through the experience of a human co-worker with a robot, this letter presents the mathematics of updating α_n and β_m of a beta probability distribution based on a reward function $r_{\psi^*} := r_{\psi^*}(s_t, a_t) \mid (s_t, a_t) \in \xi^R$ output at each timestep q . In Eq. (8), the subscripts n and m indicate the total number of successful and unsuccessful state-action counts, respectively.

$$\alpha_n = \begin{cases} \sum_{i=0}^{n-1} (\gamma^i \cdot \alpha_{n-i-1}), & \text{if } r_{\psi^*} \leq \varepsilon \\ \sum_{i=0}^{n-1} (\gamma^i \cdot \alpha_{n-i-1}) + \omega_n^s \cdot r_{\psi^*}, & \text{if } r_{\psi^*} > \varepsilon \end{cases} \quad (8)$$

$$\beta_m = \begin{cases} \sum_{j=0}^{m-1} (\gamma^j \cdot \beta_{m-j-1}), & \text{if } r_{\psi^*} > \varepsilon \\ \sum_{j=0}^{m-1} (\gamma^j \cdot \beta_{m-j-1}) + \omega_m^f \cdot e^{|r_{\psi^*}|}, & \text{if } r_{\psi^*} \leq \varepsilon \end{cases}$$

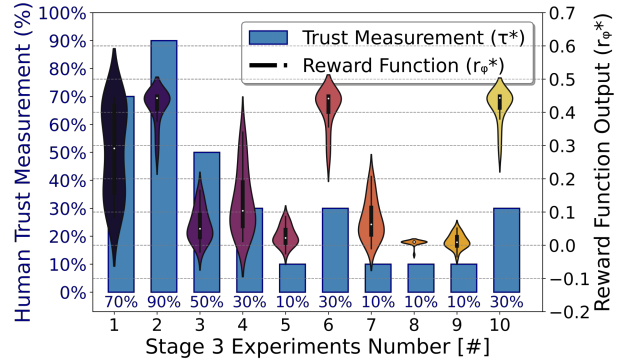


Fig. 5. **Relationship between self-reported trust and reward in Stage-3.** Violin plots illustrate the distribution of reward values. Trust is self-reported by the human using a constant scale (Fig. 3). Although experiments 6 and 10 had high rewards, trust remained lower than in experiment 2. This suggests that prior interactions influenced current trust, which underscores the importance of incorporating historical context into trust estimation.

Unlike [13], which models trust as a deterministic linear function of performance, the proposed framework uses a probabilistic approach. This allows for trust estimation at granular timescales with a more dynamic representation.

The works in [7], [36] highlight that human trust is history-dependent. To account for this characteristic of human trust, updating the α_n and β_m parameters with a weighted aging technique is chosen, as shown in Eq. (8). This technique includes a discount factor ($0 < \gamma \leq 1$), reflecting the importance of previous timestep distribution parameters (α_{n-1} , β_{m-1}). The primary reason for this interpretation is that human trust is cumulative and interaction history-dependent [26]. Furthermore, the assumption in this letter is that $r_{\psi^*} > \varepsilon$ represents success, whereas $r_{\psi^*} \leq \varepsilon$ denotes failure. Therefore, in Eq. (8), the constant number notation of the success factor is ω_n^s and the failure reward factor is ω_m^f . Overall, trust estimation depends on the parameter set $(\alpha_0, \beta_0, \omega_n^s, \omega_m^f, \varepsilon, \gamma)$ and reward function $r_{\psi^*}(s_t, a_t)$.

The mean of the beta distribution in Eq. (7) estimates human trust at a given state s_t and action a_t . Unlike [19], the proposed framework updates the trust distribution at each timestep rather than only at the end of the task. Therefore, this results in an estimation at more fine-grained timescales.

3) *Maximum Likelihood Estimation:* In Stage-3 of Fig. 2, the iterative process improves the mapping between reward values and human trust measurements τ_q^* using a maximum likelihood estimation (MLE) with a parameter set of $\lambda = \{\alpha_0, \beta_0, \omega_n^s, \omega_m^f, \varepsilon, \gamma\}$. Eq. (8) shows a threshold parameter ε in the non-differentiable piece-wise function for the updates at granular timescales. We use a popular derivative-free differential evolution method [37] to optimize λ . The objective is to minimize the negative log-likelihood between τ_q^* and $\hat{\tau}_q$ after each experiment in Stage-3.

C. Major Empirical Findings on Trust Dynamics

1) *History Dependency:* Trust at the previous timestep, τ_{q-1} , influences trust at the immediate next timestep, τ_q . Research in [26] highlights this history-dependent nature of trust. The proposed beta reputation in Eq. (8) mathematically captures this characteristic of trust with an aging factor γ .

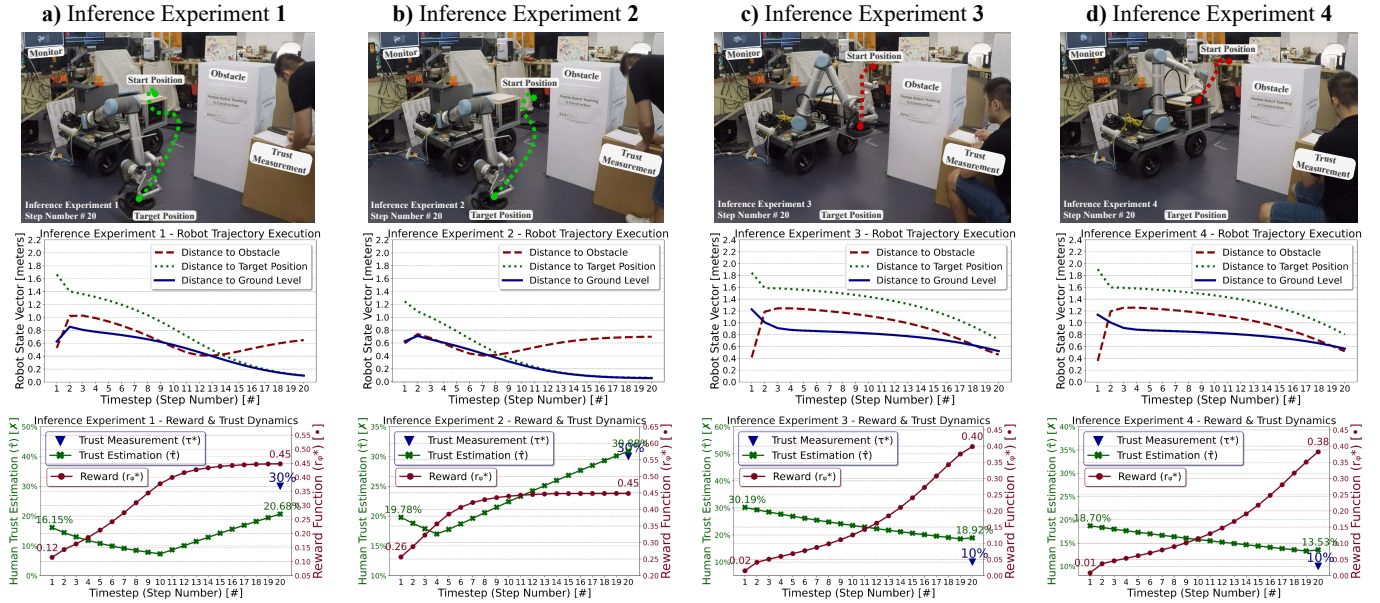


Fig. 6. **Evaluation of the proposed framework.** Each column (a-d) represents distinct and consecutive verification experiments in Stage-4 of Fig. 2. The results show four tests to evaluate human trust estimation accuracy and observe corresponding reward function values. These experiments highlight the granularity and history dependency of human trust estimation. After each experiment, the trust estimation builds on the previous experiment, reflecting the cumulative nature of human trust. In contrast, reward function values are independent of collaboration history. Experiments a-b were successful; the robot end-effector reached the target position. Experiments c-d were unsuccessful; the robot end-effector failed to reach the target position. Note that human trust measurements are discrete, but human trust estimation is at granular timescales, which dynamically change at each timestep of the experiments.

2) *Impact of Adverse Experiences:* In inference experiment 3, the robot failed to bring the tile to the target position, as shown in Fig. 6. Trust estimation at the first timestep of this experiment was 30.19%, and it continued to decrease, reaching 18.92% by the end of the experiment. For verification, human trust was self-reported as “Moderate Distrust” after inference experiment 2 and “High Distrust” at the end of inference experiment 3. This suggests that the human trust was affected by an adverse experience. This observation aligns with previous findings in the literature, which indicate that a negative experience significantly impacts trust [10].

3) *Convergence of Human Trust:* When $n, m \rightarrow \infty$, $\sum_{i=0}^{n-1} (\gamma^i \cdot \alpha_{n-i-1})$ and $\sum_{j=0}^{m-1} (\gamma^j \cdot \beta_{m-j-1})$ in Eq. (8) diminish due to a discount factor ($0 < \gamma \leq 1$). As a result, the contribution of earlier interactions gradually becomes negligible. Additionally, the terms $\omega_n^s \cdot r_{\psi^*}$ and $\omega_m^f \cdot e^{|r_{\psi^*}|}$ in Eq. (8) only shift the beta distribution by a constant factor without introducing instability in the proposed beta reputation. Consequently, because parameters θ^* and ψ^* are constant in Stage-3 and Stage-4, trust estimations converge to a stable condition after repeated collaboration with the same robot, as stated in the previous research in [28].

V. EXPERIMENTAL EVALUATION AND RESULTS

We conducted a case study to evaluate our framework for fine-grained trust estimation in a tiling task, where a robot transported a tile while avoiding an obstacle (Fig. 1). The experiment included two stages: Stage-3 with iterative model updates and Stage-4 using a fixed model. One participant (a 25-year-old male with prior technical experience in HRC) completed multiple experiments, and after each experiment, this participant self-reported trust using the scale in Fig. 3.

The state vector ($s_t \in \mathcal{S}$) included distances to the obstacle, the ground, and the target position, similar to [21], [22]. The action vector ($a_t \in \mathcal{A}$) represented the x, y, and z positions of the robot’s end-effector at each timestep t .

A. Reward Function and Human Trust Measurements

In each Stage-3 cycle (Fig. 2), a human randomly set the robot’s initial position at the start location. The robot then executed a policy $\pi_{\theta}(a_t | s_t)$ to transport a tile to a target location while avoiding an obstacle.

Trust measurements did not always align with reward values (see Fig. 5), as the reward function does not account for the history of HRC. In contrast, the beta reputation captures history-dependent trust through cumulative interactions.

With the proposed framework, the robot estimated decreasing levels of trust in the subsequent experiments of Stage-3 (see Fig. 4). This downward trend corresponds to the low reward values and associated trust measurements in Fig. 5. The distribution widens with increasing experiment number in Fig. 4, while the opposite trend appears in Fig. 7. This discrepancy arises because, during Stage-3, trust estimation parameters are updated after each experiment, whereas in Stage-4, the parameters are constant.

B. Analysis of Verification Results

In the Stage-4 experiments, the robot successfully completed the task in experiments 1 and 2 but failed in experiments 3 and 4. As shown in Fig. 6, its failure to bring a tile to the target location led to lower trust measurements. Trust estimation also declined in experiments 3 and 4, which aligned with reduced reward values. These findings align well with

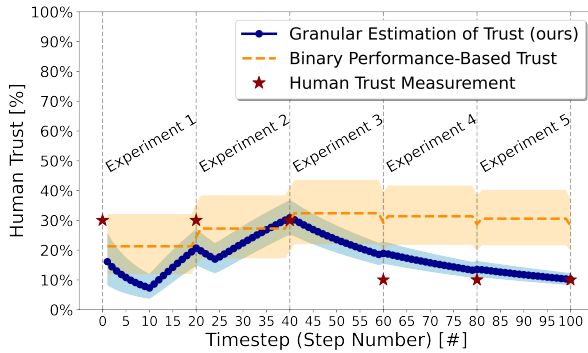


Fig. 7. **Comparison of trust estimation models.** The binary performance-based model only updates at the end of each experiment, as it relies on the overall success or failure of the task, resulting in a static trust estimation that remains unchanged throughout the experiment.

the assumption presented in Section III-C that human trust is primarily affected by the robot’s performance.

1) *History Dependency*: At the end of inference experiment 1, trust was estimated at 20.68%, and reduced to 19.78% at the start of experiment 2 (see Fig. 6). As shown in Fig. 7, the model updates trust distributions at each timestep. It reflects history-dependent continuity across all Stage-4 experiments. This aligns with the view that trust in human-machine systems should be continuously evolving [26].

2) *Impact of Adverse Experiences*: In Stage-4, the parameters λ^* , ψ^* , and θ^* were fixed. After the MLE process in Stage-3, the success and failure weights were $\omega^s = 3.7897$ and $\omega^f = 4.5390$, respectively. These results align with prior findings [10], [19] that failures impact trust more than successes ($\omega^f \geq \omega^s$), as discussed in Section IV-C).

3) *Nonlinearity in Trust Estimations*: In experiment 1, estimated trust decreased from 16.15% to below 10% by timestep 10, then increased to 20.68% by the termination timestep ($T = 20$) (see Fig. 6). This pattern results from the piecewise beta reputation function, which depends on the threshold parameter ε (Eq. (8)). A similar trend appears in experiment 2, where trust decreased until the timestep 4 due to the robot approaching the obstacle too quickly, resulting in reward values below ε . In experiments 3 and 4, estimated trust remained low because the robot failed to complete the task. These results are consistent with prior work [19], [38].

C. Comparison with Binary Performance-Based Trust Model

To evaluate the granular trust estimation model, we conducted comparative experiments in Stage-4 using the binary performance-based trust model from [19]. As the source code was unavailable, we implemented the baseline and optimized both models via MLE using differential evolution [37]. We calculated absolute error as the difference between the trust measurement and the trust estimation at the final timestep of each experiment (see Table I). A key advantage of our framework is its ability to provide trust estimations at fine-grained timescales without requiring labor-intensive, hand-crafted performance indicator design. However, a limitation is that the variance of the trust distribution tends to decrease over time, as the beta reputation model becomes more confident with each timestep. This suggests a future direction:

TABLE I
ABSOLUTE ERRORS AT THE END OF EACH INFERENCE EXPERIMENT.

Inference Number	Self-Reported Human Trust	Absolute Errors (%)			
		Binary Performance Based Model		Granular Estimation of Trust (ours)	
		μ	σ_{\max}^2	μ	σ_{\max}^2
1	Moderate Distrust	5.61	15.39	9.31	15.12
2	Moderate Distrust	0.09	11.12	0.87	6.74
3	High Distrust	19.17	29.56	8.91	12.78
4	High Distrust	18.61	28.33	3.52	6.35
5	High Distrust	18.17	27.32	0.18	2.35

Bold values indicate which model has less error in each experiment. μ is the absolute error between self-reported trust and the mean of the estimated trust distribution; σ_{\max}^2 is the maximum variance of the error.

exploring ways to regulate model confidence not only based on interaction number but also on contextual dynamics.

VI. CONCLUSION

The proposed framework estimates trust at each timestep, offering a fine-grained view of trust dynamics. It also avoids manual performance metric design by using maximum entropy optimization to generate a continuous reward function.

While this work focuses on estimating trust at fine-grained timescales, it has three main limitations. First, the framework does not yet integrate trust estimation with robot control. However, this is a promising direction for future work. For instance, prior research [39] adapted robot behavior based on trust, while other studies [9], [24] explored trust repair strategies, such as apology, denial, and transparency. Since our framework estimates trust at each timestep, it offers the potential for real-time behavioral adaptation during a task. Second, the trust measurements in this study rely on self-reports, which limits validation to discrete snapshots rather than continuous measurement at every timestep. There are ongoing works in the development of real-time measures of behavioral, physiological, and neural signatures of trust [7], [11], including electroencephalography, eye-tracking, and electrodermal activity measures [40]. Third, using a Likert scale comes with a key limitation. A Likert scale is an ordinal measure, meaning the distance between response categories may vary across points and individuals [41]. This variability can introduce biases based on individual response styles. Despite this limitation, the Likert scale remains a pragmatic and cost-effective tool. It is widely used in psychology to capture subjectivity and has demonstrated strong predictive validity [42]. An alternative to address this limitation is iterative optimization using preference-based data [33]. In one study [34], participants provided their preferences to improve the robot’s performance and thereby increase human trust. However, this approach does not allow for the estimation of trust dynamics throughout the interaction. In contrast, accurately estimating the dynamics of human trust is necessary before applying strategies to mitigate undertrust and overtrust in the future [9], [24], [39]. While this work differs from previous studies by focusing on the mathematical modeling of trust at fine-grained timescales, future research will explore

ways to incorporate individual differences in experiments involving multiple participants.

ACKNOWLEDGEMENT

The authors would like to acknowledge Gibson Hu for technical support.

REFERENCES

- [1] B. King-Casas, D. Tomlin, C. Anen, C. F. Camerer, S. R. Quartz, and P. R. Montague, "Getting to know you: reputation and trust in a two-person economic exchange," *Science*, vol. 308, no. 5718, pp. 78–83, 2005.
- [2] D. L. Ferrin, M. C. Bligh, and J. C. Kohles, "It takes two to tango: An interdependence analysis of the spiraling of perceived trustworthiness and cooperation in interpersonal and intergroup relationships," *Organizational behavior and human decision processes*, vol. 107, no. 2, pp. 161–178, 2008.
- [3] S. Lewandowsky, M. Mundy, and G. Tan, "The dynamics of trust: comparing humans to automation," *Journal of Experimental Psychology: Applied*, vol. 6, no. 2, p. 104, 2000.
- [4] M. Lewis, K. Sycara, and P. Walker, "The role of trust in human-robot interaction," *Foundations of trusted autonomy*, pp. 135–159, 2018.
- [5] A. Gunia, "The role of trust in human-machine interaction: Cognitive science perspective," in *Artificial Intelligence, Management and Trust*. Routledge, 2024, pp. 85–126.
- [6] A. Xu and G. Dudek, "Optimo: Online probabilistic trust inference model for asymmetric human-robot collaborations," in *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*, 2015, pp. 221–228.
- [7] C. S. Nam and J. B. Lyons, *Trust in human-robot interaction*. Academic Press, 2020.
- [8] P. Kaniarasu, A. Steinfeld, M. Desai, and H. Yanco, "Robot confidence and trust alignment," in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2013, pp. 155–156.
- [9] E. J. De Visser, M. M. Peeters, M. F. Jung, S. Kohn, T. H. Shaw, R. Pak, and M. A. Neerincx, "Towards a theory of longitudinal trust calibration in human-robot teams," *International journal of social robotics*, vol. 12, no. 2, pp. 459–478, 2020.
- [10] X. J. Yang, C. D. Wickens, and K. Hölttä-Otto, "How users adjust trust in automation: Contrast effect and hindsight bias," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 60, no. 1, pp. 196–200, 2016.
- [11] Y. Wang, F. Li, H. Zheng, L. Jiang, M. F. Mahani, and Z. Liao, "Human trust in robots: A survey on trust models and their controls/robotics applications," *IEEE Open Journal of Control Systems*, 2023.
- [12] M. Lewis, H. Li, and K. Sycara, "Deep learning, transparency, and trust in human robot teamwork," in *Trust in human-robot interaction*. Elsevier, 2021, pp. 321–352.
- [13] Q. Wang, D. Liu, M. G. Carmichael, S. Aldini, and C.-T. Lin, "Computational model of robot trust in human co-worker for physical human-robot collaboration," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3146–3153, 2022.
- [14] K. J. Williams, M. S. Yuh, and N. Jain, "A computational model of coupled human trust and self-confidence dynamics," *ACM transactions on human-robot interaction*, vol. 12, no. 3, pp. 1–29, 2023.
- [15] P. A. Hancock, D. R. Billings, K. E. Schaefer, J. Y. C. Chen, E. J. de Visser, and R. Parasuraman, "A meta-analysis of factors affecting trust in human-robot interaction," *Human Factors*, vol. 53, no. 5, pp. 517–527, 2011.
- [16] K. E. Schaefer, "Measuring trust in human robot interactions: Development of the "trust perception scale-hri"," in *Robust intelligence and trust in autonomous systems*. Springer, 2016, pp. 191–218.
- [17] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa, "Planning with trust for human-robot collaboration," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, 2018, p. 307–315.
- [18] C. Nam, P. Walker, H. Li, M. Lewis, and K. Sycara, "Models of trust in human control of swarms with varied levels of autonomy," *IEEE Transactions on Human-Machine Systems*, vol. 50, no. 3, pp. 194–204, 2019.
- [19] Y. Guo and X. J. Yang, "Modeling and predicting trust dynamics in human-robot teaming: A bayesian inference approach," *International Journal of Social Robotics*, vol. 13, no. 8, pp. 1899–1909, 2021.
- [20] S. Bhat, J. B. Lyons, C. Shi, and X. J. Yang, "Clustering trust dynamics in a human-robot sequential decision-making task," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8815–8822, 2022.
- [21] E. Biryk, N. Huynh, M. J. Kochenderfer, and D. Sadigh, "Active preference-based gaussian process regression for reward learning and optimization," *The International Journal of Robotics Research*, 2023.
- [22] E. Biryk, D. P. Losey, M. Palan, N. C. Landolfi, G. Shevchuk, and D. Sadigh, "Learning reward functions from diverse sources of human feedback: Optimally integrating demonstrations and preferences," *The International Journal of Robotics Research*, vol. 41, no. 1, pp. 45–67, 2022.
- [23] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," in *Proceedings of The 33rd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 48. PMLR, 2016, pp. 49–58.
- [24] Y. Li and F. Zhang, "Trust-preserved human-robot shared autonomy enabled by bayesian relational event modeling," *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 10716–10723, 2024.
- [25] B. Wu, B. Hu, and H. Lin, "Toward efficient manufacturing systems: A trust based human robot collaboration," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 1536–1541.
- [26] L. John and N. Moray, "Trust, control strategies and allocation of function in human-machine systems," *Ergonomics*, vol. 35, no. 10, pp. 1243–1270, 1992.
- [27] S. M. Fleming, "Metacognition and confidence: A review and synthesis," *Annual Review of Psychology*, vol. 75, no. 1, pp. 241–268, 2024.
- [28] X. J. Yang, V. V. Unhelkar, K. Li, and J. A. Shah, "Evaluating effects of user experience and system transparency on trust in automation," in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, 2017, p. 408–416.
- [29] A. Jøsang and R. Ismail, "The beta reputation system," in *Proceedings of the 15th bled electronic commerce conference*, 2002, pp. 324–337.
- [30] B. F. Malle, K. Fischer, J. Young, A. Moon, and E. Collins, "Trust and the discrepancy between expectations and actual capabilities," *Human-robot interaction: Control, analysis, and design*, pp. 1–23, 2020.
- [31] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human Factors*, vol. 46, no. 1, pp. 50–80, 2004.
- [32] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *Proc. AAAI*, 2008, pp. 1433–1438.
- [33] G. Campagna, M. Lagomarsino, M. Lorenzini, D. Chrysostomou, M. Rehm, and A. Ajoudani, "Promoting trust in industrial human-robot collaboration through preference-based optimization," *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 9255–9262, 2024.
- [34] A. Bemporad and D. Piga, "Global optimization based on active preference learning with radial basis functions," *Machine Learning*, vol. 110, no. 2, pp. 417–448, 2021.
- [35] G. Swamy, D. Wu, S. Choudhury, D. Bagnell, and S. Wu, "Inverse reinforcement learning without reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2023, pp. 33299–33318.
- [36] V. Frey and J. Martinez, "Interpersonal trust modelling through multi-agent reinforcement learning," *Cognitive Systems Research*, vol. 83, p. 101157, 2024.
- [37] R. Storn and K. Price, "Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces," *Journal of global optimization*, vol. 11, pp. 341–359, 1997.
- [38] H. Azevedo-Sa, S. K. Jayaraman, C. T. Esterwood, X. J. Yang, L. P. Robert Jr, and D. M. Tilbury, "Real-time estimation of drivers' trust in automated driving systems," *International Journal of Social Robotics*, vol. 13, no. 8, pp. 1911–1927, 2021.
- [39] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa, "Trust-aware decision making for human-robot collaboration: Model learning and planning," *J. Hum.-Robot Interact.*, vol. 9, no. 2, Jan. 2020.
- [40] S. C. Kohn, E. J. de Visser, E. Wiese, Y.-C. Lee, and T. H. Shaw, "Measurement of trust in automation: A narrative review and reference guide," *Frontiers in Psychology*, vol. Volume 12 - 2021, 2021.
- [41] B. Weijters, M. Geuens, and N. Schillewaert, "The stability of individual response styles," *Psychological methods*, vol. 15, no. 1, pp. 96–110, 2010.
- [42] C. L. Eddy, K. C. Herman, and W. M. Reinke, "Single-item teacher stress and coping measures: Concurrent and predictive validity and sensitivity to change," *Journal of School Psychology*, vol. 76, pp. 17–32, 2019.