Trust-Aware Decision Making for Human-Robot Collaboration: Model Learning and Planning

MIN CHEN, National University of Singapore

STEFANOS NIKOLAIDIS, University of Southern California

HAROLD SOH and DAVID HSU, National University of Singapore

SIDDHARTHA SRINIVASA, University of Washington

Trust in autonomy is essential for effective human-robot collaboration and user adoption of autonomous systems such as robot assistants. This article introduces a computational model that integrates trust into robot decision making. Specifically, we learn from data a partially observable Markov decision process (POMDP) with human trust as a latent variable. The trust-POMDP model provides a principled approach for the robot to (i) infer the trust of a human teammate through interaction, (ii) reason about the effect of its own actions on human trust, and (iii) choose actions that maximize team performance over the long term. We validated the model through human subject experiments on a table clearing task in simulation (201 participants) and with a real robot (20 participants). In our studies, the robot builds human trust by manipulating low-risk objects first. Interestingly, the robot sometimes fails intentionally to modulate human trust and achieve the best team performance. These results show that the trust-POMDP calibrates trust to improve human-robot team performance over the long term. Further, they highlight that maximizing trust alone does not always lead to the best performance.

CCS Concepts: • **Human-centered computing** → **Collaborative interaction**; • **Computing methodologies** → *Planning under uncertainty*;

Additional Key Words and Phrases: Trust models, human-robot collaboration, partially observable Markov decision process (POMDP)

ACM Reference format:

Min Chen, Stefanos Nikolaidis, Harold Soh, David Hsu, and Siddhartha Srinivasa. 2020. Trust-Aware Decision Making for Human-Robot Collaboration: Model Learning and Planning. *ACM Trans. Hum.-Robot Interact.* 9, 2, Article 9 (January 2020), 23 pages.

<https://doi.org/10.1145/3359616>

M. Chen and S. Nikolaidis contributed equally to the work. The user studies in this work were approved by the NUS-IRB: A-16-356.

This work was funded in part by the Singapore Ministry of Education (grant MOE2016-T2-2-068), the National University of Singapore (grant R-252-000-587-112), U.S. National Institutes of Health R01 (grant R01EB019335), U.S. National Science Foundation CPS (grant 1544797), U.S. National Science Foundation NRI (grant 1637748), and the Office of Naval Research. Authors' addresses: M. Chen, H. Soh, and D. Hsu, National University of Singapore, 21 Lower Kent Ridge Rd, Singapore 119077; emails: {chenmin, harold, dyhsu}@comp.nus.edu.sg; S. Nikolaidis, University of Southern California, 3710 McClintock Ave, Los Angeles, CA 90089; email: snikolai@alumni.cmu.edu; S. Srinivasa, Paul G. Allen Center, Box 352350, 185 E Stevens Way NE, Seattle, WA 98195-2350; email: siddh@cs.uw.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-9522/2020/01-ART9

<https://doi.org/10.1145/3359616>

1 INTRODUCTION

Trust is essential for seamless human-robot collaboration and user adoption of autonomous systems, such as robot assistants. Over-trusting robot autonomy may lead to misuse of such systems, where people rely excessively on automation and fail to intervene during critical failures [26]. However, a lack of trust leads to disuse of autonomous systems: users ignore the systems' capabilities, with negative effects on overall performance. We witnessed an example of users' distrust in the system in one of our studies where a human participant and a robot collaborated to clear a table (Figure 1). Although the robot was fully capable of handling all objects on the table, inexperienced participants did not trust that the robot was able to succeed and stopped the robot from moving the wine glass, since they were afraid that the glass may fall and break. It was clear that their trust was poorly calibrated with respect to the robot's true capabilities. This, in turn, had a significant effect on the interaction. This study revealed that to achieve fluent human-robot collaboration, the robot should *monitor* human trust and *influence* it so that it matches the system's capabilities. In our study, for instance, the robot should build human trust first by acting in a trustworthy manner, before going for the wine glass.

We propose a trust-based computational model of robot decision making: Since trust is not fully observable, we model it as a latent variable in a partially observable Markov decision process (POMDP) [22]. Our trust-POMDP model contains two key components: (i) a trust dynamics model, which captures the evolution of human trust in the robot, and (ii) a human decision model, which connects trust with human actions. Our POMDP formulation can accommodate a variety of trust dynamics and human decision models. Here, we adopt a data-driven approach and learn these models from data.

Although prior work has studied human trust elicitation and modeling [14, 25, 44, 46], we close the loop between trust modeling and robot decision making. The trust-POMDP enables the robot to systematically infer and influence the human collaborator's trust, and leverage trust for improved human-robot collaboration and long-term task performance.

Consider again the table clearing example (Figure 2). The trust-POMDP strategy first removes the three plastic water bottles to build up trust and only attempts to remove the wine glass afterward. In contrast, a baseline myopic strategy maximizes short-term task performance and does not account for human trust in choosing the robot actions. It first removes the wine glass, which offers the highest reward, resulting in unnecessary interventions by human collaborators with low initial trust.

We validated the trust-POMDP model through human subject experiments on the collaborative table clearing task, both online in simulation (201 participants) and with a real robot (20 participants). Compared with the myopic strategy, the trust-POMDP strategy significantly reduced participants' intervention rate, indicating improved team collaboration and task performance.

In these experiments, the robot always succeeded. Robots, however, fail frequently. What if the robot is likely to *fail* when picking up the wine glass? The robot should then assess human trust in the beginning of the task; if trust is too high, the robot should effectively *communicate* this to the human to calibrate human trust to the appropriate level. Although human teammates are able to use natural language to communicate expectations [28], our assistive robotic arm does not have verbal communication capabilities. The trust-POMDP strategy in this case enables the robot to modulate human trust by *intentionally failing* when picking up the bottles, before attempting to grasp the wine glass. This prompts the human to intervene when the robot attempts to pick up the wine glass, preventing failure.

This article builds upon our previous work [7] by introducing robot failures into the computational framework. In particular, (i) we augment the dynamics model with robot failures, add a new session of data collection to learn the model, and discuss the effect of failures on different levels of

Fig. 1. A robot and a human collaborate to clear a table. The human, with low initial trust in the robot, intervenes to stop the robot from moving the wine glass.

Fig. 2. Sample runs of the trust-POMDP strategy and the myopic strategy on a collaborative table clearing task. The top row shows the probabilistic estimates of human trust over time on a 7-point Likert scale. The trust-POMDP strategy starts by moving the plastic bottles to build trust ($T = 1, 2, 3$) and moves the wine glass only when the estimated trust is high enough ($T = 5$). The myopic strategy does not account for trust and starts with the wine glass, causing the human, who has low trust initially, to intervene ($T = 1$).

trust; (ii) we simulate and visualize robot policies with the learned model; and (iii) we provide an analysis of the results in the case of an adaptive policy that enables the robot to assess participants' initial trust and intentionally fail.

Integrating trust modeling and robot decision making enables robot behaviors that leverage human trust and actively *modulate* it for seamless human-robot collaboration. Under the trust-POMDP model, the robot deliberately chooses to fail to reduce the trust of an overly trusting user and achieve better task performance over the long term. Further, embedding trust in a reward-based POMDP framework makes our robot task driven: When human collaboration is unnecessary, the robot may set aside trust building and act to maximize the team task performance directly. All of these diverse behaviors emerge automatically from the trust-POMDP model, without explicit manual robot programming.

2 RELATED WORK

Trust is a complex issue and has been studied extensively in the social science research literature [16, 23], with Mayer et al. [29] suggesting that three general levels summarize the bases of trust: ability, integrity, and benevolence. Trust in automation differs from trust between people in that (i) humans are more collaborative and trusting of human partners even when observed agent actions are identical [45], and (ii) automation lacks intentionality [26]. Additionally, in a human-robot collaboration task, the human and robot share the same objective metric of task performance. Therefore, similar to previous work [9, 35, 36, 44, 47], we focus on one particular dimension of trust and assume that human teammates will not expect the robot to deceive them on purpose, and their trust will depend mainly on the *perceived robot ability* to complete the task successfully.

Binary measures of trust [17], as well as continuous measures [9, 25, 47], and ordinal scales [18, 20, 30] have been proposed. For real-time measurement, Desai [9] proposed the Area Under Trust Curve (AUTC) measure, which was recently used to account for one's entire interactive experience with the robot [48].

Researchers have also studied the temporal dynamics of trust conditioned on the task performance: Lee and Moray [25] proposed an autoregressive moving average vector form of time series analysis; Floyd et al. [14] used case-based reasoning; Xu and Dudek [46] proposed an online probabilistic trust inference model to estimate a robot's trustworthiness; Wang et al. [44] showed that adding transparency in the robot model by generating explanations improved trust and performance in human teams; Desai et al. [10, 11] showed that robot failures had a negative impact on the human trust, and early robot failures led to dramatically lower trust than later robot failures. Previous works have focused on either quantifying trust or studying the dynamics of trust in human-robot interaction; however, our work enables the robot to leverage upon a model of human trust and choose actions to maximize task performance.

In human-robot collaborative tasks, the robot often needs to reason over the human's hidden mental state in its decision making. The POMDP provides a principled general framework for such reasoning. It has enabled robotic teammates to coordinate through communication [5] and software agents to infer the intention of human players in game AI applications [27]. The model has been successfully applied to real-world tasks, such as autonomous driving where the robot car interacts with pedestrians and human drivers [2, 3, 15]. When the state and action space of the POMDP model become continuous, one can use hindsight optimization [19], or value of information heuristics [39], which generate approximate solutions but are computationally more efficient.

Nikolaidis et al. [33] proposed to infer the human type or preference online using models learned from joint-action demonstrations. This formalism recently extended from one-way adaptation (from robot to human) to human-robot *mutual* adaptation [31, 32], where the human may choose to change his or her preference and follow a policy demonstrated by the robot in the recent

history. In this work, we provide a general way to link the whole interaction history with the human policy, by incorporating human trust dynamics into the planning framework.

3 THE TRUST-POMDP

3.1 Human-Robot Team Model

We formalize the human-robot team as a Markov decision process (MDP), with world state $x \in X$, robot action $a^R \in A^R$, and human action $a^H \in A^H$. The system evolves according to a probabilistic state transition function $p(x'|x, a^R, a^H)$ that specifies the probability of transitioning from state x to state x' when actions a^R and a^H are applied in state x . After transitioning, the team receives a real-valued reward $r(x, a^R, a^H)$, which is constructed to elicit the desirable team behaviors.

We denote by $h_t = \{x_0, a_0^R, a_0^H, x_1, r_1, \dots, x_{t-1}, a_{t-1}^R, a_{t-1}^H, x_t, r_t\} \in H_t$ as the history of interaction between robot and human until timestep t . In this work, we assume that the human observes the robot's current action and then decides his or her own action. In the most general setting, the human uses the entire interaction history h_t to decide the action. Thus, we can write the human's (possibly stochastic) policy as $\pi^H(a_t^H|x_t, a_t^R, h_t)$, which outputs the probability of each human action a_t^H .

Given a robot policy π^R , the *value* (i.e., the expected total discounted reward of starting at a state x_0 and following the robot and human policies) is

$$v(x_0|\pi^R, \pi^H) = \mathbb{E}_{a_t^R \sim \pi^R, a_t^H \sim \pi^H} \sum_{t=0}^{\infty} \gamma^t r(x_t, a_t^R, a_t^H), \quad (1)$$

where x_t is the state at timestep t , and a_t^R and a_t^H are the robot action and human action at timestep t . γ is the discount factor that favors immediate rewards over future ones.

The robot's optimal policy π_*^R can be computed as

$$\pi_*^R = \arg \max_{\pi^R} v(x_0|\pi^R, \pi^H). \quad (2)$$

In our case, however, the robot does not know the human policy in advance. It computes the optimal policy under expectation over the human policy:

$$\pi_*^R = \arg \max_{\pi^R} \mathbb{E}_{\pi^H} v(x_0|\pi^R, \pi^H). \quad (3)$$

Key to solving Equation (3) is for the robot to model the human policy, which potentially depends on the entire history h_t . The history h_t may grow arbitrarily long and make the optimization extremely difficult.

3.2 Trust-Dependent Human Behaviors

Our insight is that in a number of human-robot collaboration scenarios, trust is a compact approximation of the interaction history h_t . This allows us to condition human behavior on the inferred trust level and in turn find the optimal policy that maximizes team performance. Following previous work on trust modeling [46], we assume that trust can be modeled as a single scalar random variable θ . Thus, the human policy is approximated as

$$\pi^H(a_t^H|x_t, a_t^R, \theta_t) \sim \pi^H(a_t^H|x_t, a_t^R, h_t). \quad (4)$$

3.3 Trust Dynamics

Human trust changes over time. We adopt a common assumption on the trust dynamics: trust evolves based on the robot's performance e_t [25, 46]. Performance can depend not just on the

Fig. 3. The trust-POMDP graphical model (left) and the team interaction flowchart (right). The human action a_t^H depends on the world state x_t and the human trust θ_t . The robot action a_t^R depends on the world state x_t and its belief over trust θ_t . The world state x_t transits to x_{t+1} when provided the human action a_t^H and the robot action a_t^R . The human trust θ_t transits to θ_{t+1} when provided the robot performance e_{t+1} .

current and transitioned world state but also the human and robot's actions:

$$e_{t+1} = \text{performance}(x_{t+1}, x_t, a_t^R, a_t^H). \quad (5)$$

For example, performance may indicate success or failure of the robot to accomplish a task. This allows us to write our trust dynamics equation as

$$\theta_{t+1} \sim p(\theta_{t+1} | \theta_t, e_{t+1}). \quad (6)$$

We detail in Section 4 how trust dynamics is learned via interaction.

3.4 Maximizing Team Performance

Trust cannot be directly observed by the robot and therefore must be inferred from the human's actions. In addition, armed with a model, the robot may actively modulate the human's trust for the team's long-term reward.

We achieve this behavior by modeling the interaction as a POMDP, which provides a principled general framework for sequential decision making under uncertainty. A graphical model of the trust-POMDP and a flowchart of the interaction are shown in Figure 3.

To build a trust-POMDP, we augment the state space to include the partially observed human trust θ —for instance, the state $s = (x, \theta)$ is composed of the fully observed world state x and the partially observed human trust θ . We maintain a belief b over the human's trust. The trust dynamics and human behavioral policy are embedded in the transition dynamics of trust-POMDP. We describe in Section 4 how we learn the trust dynamics and the human behavioral policy.

The robot now has two distinct objectives through its actions:

- *Exploitation*: Maximize the team's reward.
- *Exploration*: Reveal and change the human's trust so that future actions are rewarded better.

The solution to a trust-POMDP is a policy that maps belief states to robot actions, (i.e., $a^R = \pi^R(b_t, x_t)$). To compute the optimal policy, we use the SARSOP algorithm [24], which is computationally efficient and has been used in various robotic tasks [3].

Table 1. Reward Function R for the Table Clearing Task

	Bottle	Fish Can	Wine Glass
SP-success	1	2	3
SP-fail	0	-4	-9
IT	0	0	0

4 LEARNING TRUST DYNAMICS AND HUMAN BEHAVIORAL POLICIES

Nested within the trust-POMDP is a model of human trust dynamics $p(\theta_{t+1}|\theta_t, e_{t+1})$ and a model of human behaviors $\pi^H(a_t^H|x_t, a_t^R, \theta_t)$. These two models are unknown to us and need to be built before we can solve the trust-POMDP and plan a robot policy. We adopted a data-driven approach and built these two models for the table clearing task from the data collected in the Amazon Mechanical Turk (AMT) study. Suitable probabilistic models derived via alternative approaches can be substituted for these learned models (e.g., explicit trust modeling from the social science field).

4.1 Data Collection

Table clearing task. A human and a robot collaborate to clear objects off a table. The objects include three water bottles, one fish can, and one wine glass. At each timestep, the robot picks up one of the remaining objects. Once the robot starts moving toward the intended object, the human can choose between two actions: {intervene and pick up the object that the robot is moving toward, stay put and let the robot pick the object by itself}. This process is repeated until all objects are cleared from the table.

Each object is associated with a different reward, based on whether the robot successfully clears it from the table (which we call *SP-success*), the robot fails in clearing it (*SP-fail*), or the human intervenes and puts it on the tray (*IT*). Table 1 shows the rewards for each object and outcome. We assume that a robot success is always better than a human intervention, as it reduces human effort. Additionally, there is no penalty if the robot fails by dropping one of the sealed water bottles, since the human can pick it up. However, dropping the fish can result in some penalty, because its contents will be spilled on the floor. Breaking the glass results in the highest penalty. We see that staying put when the robot attempts to pick up the bottle has the lowest risk, since there is no penalty if the robot fails. Yet staying put in the case of the glass object has the largest risk-return tradeoff. We expect the human to let the robot pick up the bottle even if the human's trust is low, as there is no penalty if the robot fails. But if the human does not trust the robot, we expect him or her to likely intervene on the glass or can rather than risking a high penalty in case of robot failure.

Participants. For the data collection, we recruited 231 participants through AMT.¹ The participants were all from the United States, aged 18 to 65 years and with an approval rate higher than 95%. Each participant was compensated \$1 for completing the study. To ensure the quality of the recorded data, we asked all participants an attention check question that tested their attention to the task. We removed 9 data points either because the participants failed on the attention check question or the their data were incomplete. This left us 222 valid data points for model learning.

Procedure. Each participant was asked to perform an online table clearing task together with a robot. Before the task started, the participant was informed of the reward function in Table 1. We

¹We conducted two sessions of data collection: one where the robot always succeeded and one when the robot failed with high probability. Our previous work [7] presented the results of the first session only.

Table 2. Muir's Questionnaire

-
1. To what extent can the robot's behavior be predicted from moment to moment?
 2. To what extent can you count on the robot to do its job?
 3. What degree of faith do you have that the robot will be able to cope with similar situations in the future?
 4. Overall, how much do you trust the robot?
-

first collected the participant's initial trust in the robot. We used Muir's questionnaire [30], with a 7-point Likert scale as a human trust metric (i.e., trust ranges from 1 to 7). Muir's questionnaire is a well-established measure of trust in the literature [9]. It integrates the models of Barber [4] and Rempel et al. [37] to include the dimensions of persistence, technical competence, reliability, and predictability of the robot, which are important when measuring performance-centric trust (i.e., the type of trust studied in this work). The set of questions we used is listed in Table 2.

At each timestep, the participants watched a video of the robot attempting to pick up an object and were asked to choose to intervene or stay put. They then watched a video of either the robot picking up the object or them intervening based on their action selection. Then, they reported their updated trust in the robot.

We were interested in learning the trust dynamics and the human behavioral policies for any state and robot action. However, the number of open-loop² robot policies is $O(K!)$, where K is the number of objects on the table. To focus the learning on a few interesting robot policies (i.e., picking up the glass in the beginning vs. in the end), while still covering a large space of policies, we split the data collection process so that in one half of the trials the robot randomly chose a policy out of a set of pre-specified policies, whereas in the other half the robot followed a random policy.

Data format. The data we collected from each participant has the following format:

$$d_i = \{\theta_0^M, a_0^R, a_0^H, e_1, \theta_1^M, \dots, a_{K-1}^R, a_{K-1}^H, e_K, \theta_K^M\},$$

where K is the number of objects on the table. θ_t^M is the estimated human trust at time t by averaging the participants' responses to the Muir's questionnaire to a single rating between 1 and 7. a_t^R is the action taken by the robot at timestep t . a_t^H is the action taken by the human at timestep t . e_{t+1} is the performance of the robot that indicates whether the robot succeeded at picking up the object, the robot failed, or the human intervened.

4.2 Trust Dynamics Model

We model human trust evolution as a linear Gaussian system, which is similar to Xu and Dudek [46]. Our trust dynamics model relates the human trust causally to the robot task performance e_{t+1} :

$$\begin{aligned} P(\theta_{t+1} | \theta_t, e_{t+1}) &= \mathcal{N}(\alpha_{e_{t+1}} \theta_t + \beta_{e_{t+1}}, \sigma_{e_{t+1}}^2) \\ \theta_t^M &\sim \mathcal{N}(\theta_t, \sigma^2), \quad \theta_{t+1}^M \sim \mathcal{N}(\theta_{t+1}, \sigma^2), \end{aligned} \tag{7}$$

where $\mathcal{N}(\mu, \sigma)$ denotes a Gaussian distribution with mean μ and standard deviation σ . $\alpha_{e_{t+1}}$ and $\beta_{e_{t+1}}$ are linear coefficients for the trust dynamics, given the robot task performance e_{t+1} . In the table clearing task, e_{t+1} indicates whether the robot succeeded at picking up an object, the robot

²When collecting data from AMT, the robot follows an open-loop policy (i.e., it does not adapt to the human behavior).

Fig. 4. Trust transition matrices, which represent the change of trust given the robot performance, shown by the linearly regressed line (yellow) contrasted with the X-Y line (blue). In general, trust stays constant or decreases slightly when the human intervenes (top row). It increases when the human stays put and the robot succeeds (middle row), whereas it decreases when the robot fails (bottom row).

failed, or the human intervened. For example, e_{t+1} can represent that the robot succeeded at picking a water bottle, or that the human intervened at the wine glass. θ_t^M and θ_{t+1}^M are the observed human trust (Muir's questionnaire) at timestep t and timestep $t+1$.

The unknown parameters in the trust dynamics model include $\alpha_{e_{t+1}}$, $\beta_{e_{t+1}}$, $\sigma_{e_{t+1}}$, and σ . We performed full Bayesian inference on the model through Hamiltonian Monte Carlo sampling using the Stan probabilistic programming platform [6].

Trust in our model is a discrete scalar that ranges from 1 to 7. We computed the trust transition dynamics by sampling from the learned linear Gaussian model. For example, suppose the current trust is $\theta_t = 1$, the next trust θ_{t+1} follows a Gaussian distribution $\mathcal{N}(\alpha_{e_{t+1}}\theta_t + \beta_{e_{t+1}}, \sigma_{e_{t+1}}^2)$, where we can compute the probability of $\theta_{t+1} = 1, \theta_{t+1} = 2, \dots, \theta_{t+1} = 7$, and we normalize the probability afterward. Figure 4 shows the trust transition matrices for all trust levels and all possible robot performance in the table clearing task. As we can see, human trust in the robot gradually increased with observations of successful robot actions (as indicated by transitions to higher trust levels when the participants stayed put and robot succeeded), whereas it decreased with observations of robot failures. Trust tended to remain constant or decrease slightly when interventions occurred. It also appears that the higher the trust, the greater the loss upon failure, and vice versa upon success. These results matched our expectations that successful robot performance positively influenced trust, whereas robot failures negatively affected trust.

4.3 Human Behavioral Policies

Our key intuition in the human model is that the human's behavior depends on the trust in the robot. To support our intuition, we consider two types of human behavioral models. The first model is a trust-free human behavioral model that ignores human trust, whereas the second is a trust-based human behavioral model that explicitly models human trust. In both human models, we assume humans follow the *softmax rule*³ when they make decisions in an uncertain environment [8]. More explicitly,

- *Trust-free human behavioral model*: At each timestep, the human selects an action probabilistically based on the actions' relative expected values. The expected value of an action depends on the human's belief on the robot to succeed and the risk of letting the robot do the task. In the trust-free human model, the human's belief on the robot success on a particular task does not change over time.
- *Trust-based human behavioral model*: Similar to the preceding model, the human follows the *softmax rule* at each timestep. However, the trust-based human model assumes that the human's belief on the robot success changes over time, and it depends on the human's trust in the robot.

Before we introduce the models, we define some notations. Let j denote the object that the robot tries to pick at timestep t . Let r_j^S be the reward if the human stays put and the robot succeeds, and let r_j^F be the reward if the human stays put and the robot fails. Let θ_t be the human trust in the robot at timestep t . $\mathcal{S}(x) = \frac{1}{1+e^{-x}}$ is the sigmoid function, which is equivalent to the softmax function in the case of binary human actions. $\mathcal{B}(p)$ is the Bernoulli distribution that takes action stay put with probability p .

The trust-free human behavioral model is as follows:

$$\begin{aligned} P_t &= \mathcal{S}(b_j r_j^S + (1 - b_j) r_j^F) \\ a_t^H &\sim \mathcal{B}(P_t), \end{aligned} \quad (8)$$

where b_j is the human's belief on the robot successfully picking up object j , and it remains constant. $0 < P_t < 1$ is the probability that the human stays put at timestep t . a_t^H is the action the human took at timestep t .

Next, we introduce the trust-based human behavioral model:

$$\begin{aligned} b_j^t &= \mathcal{S}(\gamma_j \theta_t + \eta_j) \\ P_t &= \mathcal{S}(b_j^t r_j^S + (1 - b_j^t) r_j^F) \\ \theta_t^M &\sim \mathcal{N}(\theta_t, \sigma^2), \quad a_t^H \sim \mathcal{B}(P_t), \end{aligned} \quad (9)$$

where b_j^t is the human's belief on robot success on object j at timestep t , and it depends on the human's trust in the robot. γ_j and η_j are the linear coefficients that relate the reported trust to the probability of robot succeeding at object j . $0 < P_t < 1$ is the probability that the human stays put at timestep t . θ_t^M is the observed human trust from Muir's questionnaire at timestep t , and we assume it follow a Gaussian distribution with mean θ_t and standard deviation σ . a_t^H is the action the human took at timestep t .

The unknown parameters here include b_j in the trust-free human model and γ_j, η_j, σ in the trust-based human model. We performed Bayesian inference on the preceding two models using Hamiltonian Monte Carlo sampling [6]. The trust-based human model (log-likelihood = -153.37)

³According to the *softmax rule*, the human's decision of which action to take is determined probabilistically on the actions' relative expected values.

Fig. 5. The model prediction on the mean human intervention rate with respect to trust (lef column) and the observed human intervention rate during the study (right column). Under the trust-free human behavioral model, the human intervention rate stays constant. Under the trust-based human behavioral model, the intervention rate decreases with increasing trust, which is similar to the observed human intervention rate, and the rate of decrease depends on the object (it is more sensitive to the riskier objects).

fit the collected AMT data better than the trust-free human model (log-likelihood = -156.40). The log-likelihood values are relatively low in both of those models due to the large variance among different users. Nevertheless, this result supports our notion that the prediction on the human behavior is improved when we explicitly model human trust.

Figure 5 shows the model prediction on the mean probability of human interventions with respect to trust (left column) and observed human behaviors during the study (right column). Overall, the prediction of the trust-based model is closer to the observed human behaviors during the study. For both the trust-based model and trust-free model, the human tends to intervene more on objects with higher risk—for instance, the human intervention rate on the glass is higher than on the can, which is again higher than on the bottle. The trust-free human behavioral model ignores human trust, and thus the human intervention rate does not change. However, the trust-based human behavioral model has a general falling trend, which indicates that participants are less likely to intervene when their trust in the robot is high. This is observed particularly for the highest-risk object (glass), where the object intervention rate drops significantly when the human trust score is maximum.

To summarize, the results of Sections 4.2 and 4.3 indicate the following:

- *Human trust is affected by robot performance:* Human trust can be built up by successfully picking up objects (Figure 4). In addition, it is a good strategy for the robot to start with low-risk objects (bottle), as the human is less likely to intervene even if the trust in the robot is low (Figure 5).
- *Human trust affects human behaviors:* The intervention rate on the high-risk objects could be reduced by building up human trust (Figure 5).

5 HUMAN SUBJECT EXPERIMENTS

We conducted two human subjects experiments: one on AMT with human participants interacting with recorded videos and one in our lab with human participants interacting with a real robot. The purpose of our study was to test whether the trust-POMDP robot policy would result in better *team performance* than a policy that did not account for human trust. To simplify the analysis of the different behaviors in these experiments, we had the robot *always succeed* when attempting to pick up the objects.

We had two experimental conditions, which we refer to as “trust-POMDP” and “myopic”:

- In the trust-POMDP condition, the robot uses human trust as a means to optimize the long-term team performance. It follows the policy computed from the trust-POMDP described in Section 3.4, where the robot’s perceived human policy is modeled via the trust-based human behavioral model described in Section 4.3.
- In the myopic condition, the robot ignores human trust. It follows a myopic policy by optimizing Equation (3), where the robot’s perceived human policy is modeled via the trust-free human behavioral model described in Section 4.3.

5.1 Online AMT Experiment

Hypothesis 1. *In the online experiment, the performance of teams in the trust-POMDP condition will be better than that of the teams in the myopic condition.*

We evaluated team performance by the accumulated reward over the task. We expected the trust-POMDP robot to reason over the probability of human interventions and act so as to minimize the intervention rate for the highest reward objects. The robot would do so by actively building up human trust before it goes for high-risk objects. On the contrary, the myopic robot policy was agnostic to how the human policy may change from the robot and human actions.

Procedure. The procedure is similar to the one for data collection (Section 4.1), with the difference that, rather than executing random sequences, the robot executes the policy associated with each condition. Although we kept Muir’s questionnaire in the experiment as a ground-truth measure of trust, the robot *did not use the score* but estimated trust solely from the trust dynamics model as described in Section 4.2.

Model parameters. In the formulation of Section 3.4, the observable state variable x represents the state of each object (on the table or removed). We assume a discrete set of values of trust $\theta : \{1, 2, 3, 4, 5, 6, 7\}$. The transition function incorporates the learned trust dynamics and human behavioral policies, as described in Section 4. The reward function R is given by Table 1. We used a discount factor of $\gamma = 0.99$, which favors immediate rewards over future rewards.

Subject allocation. We chose a between-subjects design to not bias the users with policies from previous conditions. We recruited 208 participants through AMT, aged 18 to 65 years and with an approval rate higher than 95%. Each participant was compensated \$1 for completing the study. We removed 7 wrong (participants failed on the attention check question) or incomplete data points. In the end, we had 101 data points for the trust-POMDP condition and 100 data points for the myopic condition.

5.2 Real-Robot Experiment

In the real-robot experiment, we followed the same robot policies, model parameters, and procedures as the online AMT experiment, only in this case the participants interacted with a real-robot in person. At each timestep, the participant could choose to stay put and let the robot pick up the object, or intervene by first pressing a button to stop the robot and then removing the object by him- or herself.

Hypothesis 2. *In the real-robot experiment, the performance of teams in the trust-POMDP condition will be better than that of the teams in the myopic condition.*

Subject allocation. We recruited 20 participants from our university, aged 21 to 65 years. Each participant was compensated \$10 for completing the study. All data points were kept for analysis: 10 data points for the trust-POMDP condition and 10 data points for the myopic condition.

Fig. 6. Comparison of the trust-POMDP and the myopic policies in the AMT experiment and the real-robot experiment.

5.3 Team Performance

We used the accumulated rewards as a measure of team performance. We performed a one-way ANOVA test of the accumulated rewards in the trust-based condition and the myopic condition. In addition, we computed the effect size (Cohen's d) for comparison between the means of accumulated rewards in the two conditions.

In the online AMT experiment, the accumulated rewards of trust-based condition was significantly larger than the myopic condition ($F(1, 199) = 7.81, p = 0.006$). The effect size $d = 0.4$ indicates a medium difference between the two means. This result supports Hypothesis 1.

Similarly, in the real-robot experiment, the accumulated rewards of the trust-based condition was significantly larger than the myopic condition ($F(1, 18) = 11.22, p = 0.004$). The effect size $d = 1.5$ indicates a huge difference between the two means. This result supports Hypothesis 2. The effect size showed that the difference in performance was more significant in the real-robot experiment compared with the online AMT experiment. This was mainly because (i) some online participants in the AMT study did not pay full attention to the experiment, and (ii) the participants in the AMT study interacted with a video while the participants in the real-robot study interacted with a real robot.

Overall, for both the online AMT experiment and the real-robot experiment, the trust-POMDP robot outperformed the myopic robot. The difference in performance occurred because participants' intervention rate in the trust-POMDP condition was significantly lower than the myopic condition (Figure 6, left column). In the online AMT experiment, the intervention rate in the trust-POMDP condition was 54% and 31% lower in the can and glass object. In the real-robot experiment, the intervention rate in the trust-POMDP condition dropped to zero (100% lower) in the can object and 71% lower in the glass object.

In the myopic condition, the robot picked the objects in the order of highest to lowest reward (glass, can, bottle, bottle, bottle). In contrast, the trust-based human behavior model influenced the trust-POMDP robot policy by capturing the fact that interventions on high-risk objects were more likely if trust in the robot was insufficient. Therefore, the trust-POMDP robot reasoned that it was

Fig. 7. Time-varying trust dynamics, where the mean of trust change is indicated by the green triangle. The same outcome has greater effect on trust when it occurs earlier than later, which is easy to see in the real-robot experiment, especially for high-risk objects (the can and glass).

better to start with the low-risk objects (bottles), build human trust (Figure 6, center column), and go for the high-risk object (glass) last. In this way, the trust-POMDP robot minimized the human intervention ratio on the glass and can object, which significantly improved team performance.

In summary, the trust-POMDP robot was able to make good decisions on whether to pick up the low-risk object to increase human trust or go directly to the high-risk object when trust is high enough. This is one main advantage that the trust-POMDP robot has over the myopic robot.

5.4 Trust Evolution

Figure 6 (center column) shows the participants' trust evolution in the online AMT experiment and the real-robot experiment. We make two key observations. First, successfully completing a task increased participants' trust in the robot. This is consistent with the human trust dynamics model we learned in Section 4.2. Second, the trust evolves similarly in the two experiment conditions (Figure 6, center column), given that fewer human interventions occurred under the trust-POMDP policy. This can be partially explained by a combination of averaging and nonlinear trust dynamics, specifically that robot performance in the earlier part of the task has a more pronounced impact on trust [9]. This is a specific manifestation of the “primacy effect,” a cognitive bias that results in a subject crediting a performer more if the performer succeeds earlier in time [21]. Figure 7 shows this time-varying aspect of trust dynamics in our experiment; the change in the mean of trust was larger if the robot succeeded earlier, most clearly seen for the can and glass objects in the real-robot experiment. As such, in the myopic condition, although there were more interventions on the glass/can at the beginning, this was averaged out by a larger increase in human trust.

Importantly, the trust-POMDP uses trust as the means to maximize task performance and builds up human trust only when necessary. For instance, the trust-POMDP robot does not maximize

human trust. In addition, it may choose to ignore or even reduce human trust for maximized performance (see Section 7 for examples).

5.5 Human Behavioral Policy

Figure 6 (right column) shows the observed human behaviors given different trust levels. In general, the trend of human intervention is consistent with the trust-based human behavioral model (Section 4.3), and participants were less likely to intervene as their trust in the robot increased. The human's action also depended on the type of object. For low-risk objects (bottles), participants allowed the robot's attempt to complete the task even if their trust in the robot was low. However, for a high-risk object (glass), participants intervened unless they trusted the robot more.

6 TIME-VARYING TRUST DYNAMICS

We observed in the experiment that robot performance in the earlier part of the task has a more pronounced impact on the human trust (i.e., trust dynamics varies over time). This is different from the trust dynamics model in Section 4.2, which is time invariant. We adapt the trust-POMDP model to incorporate this fact. The resulting robot policy is more aggressive at building human trust, especially at the beginning of a task.

6.1 Time-Varying Trust Dynamics Model

Similar to Section 4.2, we model trust evolution as a linear Gaussian system that relates trust causally to the robot task performance e_{t+1} . The key difference here is that we further condition trust evolution on the time. For the table clearing task, we manually divided the time into *earlier* (first and second step) and *later* (third through fifth steps). We formally write our time-varying trust dynamics model as follows:

$$\begin{aligned} P(\theta_{t+1} | \theta_t, m_t, e_{t+1}) &= \mathcal{N}(\alpha_{e_{t+1}, m_t} \theta_t + \beta_{e_{t+1}, m_t}, \sigma_{e_{t+1}, m_t}^2) \\ \theta_t^M &\sim \mathcal{N}(\theta_t, \sigma^2), \theta_{t+1}^M \sim \mathcal{N}(\theta_{t+1}, \sigma^2) \end{aligned} \quad (10)$$

where m_t is a binary variable that indicates the current progress of the task—for instance, $m_t = \text{earlier}$ when $t \leq 2$, and $m_t = \text{later}$ when $t > 2$. $\mathcal{N}(\mu, \sigma)$ denotes a Gaussian distribution with mean μ and standard deviation σ . α_{e_{t+1}, m_t} and β_{e_{t+1}, m_t} are the linear coefficients for the trust dynamics that explicitly conditioned on the robot performance e_{t+1} and current progress m_t . θ_t^M and θ_{t+1}^M are the observed human trust (Muir's questionnaire) at timestep t and timestep $t + 1$.

The unknown parameters in the time-varying trust dynamics model include α_{e_{t+1}, m_t} , β_{e_{t+1}, m_t} , σ_{e_{t+1}, m_t} , and σ . We used the trust data collected from the real-robot experiment, and we performed full Bayesian inference on the model through Hamiltonian Monte Carlo sampling using the Stan probabilistic programming platform [6].

Figure 8 shows the trust transition matrices at different times during the task (i.e., earlier or later), given that the robot succeeded at that step. As we can see, human trust in the robot increased more significantly when the robot succeeded earlier (top row) compared with when the robot succeeded later (bottom row).

6.2 Robot Policy with Time-Varying Trust Dynamics

We incorporate the time-varying trust dynamics model into a trust-POMDP, where we simply replace the trust dynamics model in Section 4.2 with the time-varying trust dynamics model. The resulting robot policy is shown in Figure 9, which is different from the robot policy with the time-invariant trust dynamics model (Figure 2, top row). In particular, the robot moves the can object at the second step instead of at the fourth step (Figure 2, top row). This is primarily due to the fact

Fig. 8. Time-varying trust transition matrices for the robot succeeding earlier (top row) and the robot succeeding later (bottom row). The yellow line is the linearly regressed line of the reported human trust. The blue line is the X-Y line (i.e., “trust before” = “trust after”). Overall, human trust increased more when the robot succeeded earlier in time.

Fig. 9. Sample run of the trust-POMDP strategy with time-varying trust dynamics. The robot goes for the can object at the second step, because succeeding at the can object earlier increases trust significantly more compared with later success.

that the robot moving the can object earlier increases trust more than moving the can object later (Figure 8, central column).

7 ROBOT FAILURES

The previous experimental results show that the trust-POMDP policy significantly outperforms the myopic policy that ignores trust in robot decision making, when the robot always succeeded. However, in the real world, the robot is also likely to fail, and we want to explore the behavior of the trust-POMDP when the robot may fail in its attempt to pick up an object with some known probability.

By explicitly modeling robot failures in the trust-POMDP, the robot is able to actively modulate human trust to maximize the performance. For example, the robot may intentionally fail to reduce

Fig. 10. Sample run of the trust-POMDP strategy when the robot may fail on the glass cup with high probability and the robot *cannot* intentionally fail on an action. The robot starts with the glass cup as the initial human trust is low and the human tends to intervene, which prevents robot failures.

human trust, and thus the human is more likely to intervene in the subsequent actions. In addition, we have not observed similar robot behaviors by simply maximizing human trust (i.e., maximizing human trust alone may not maximize the team performance).

In our experiment, we *manually* set the parameters of our robot to make it fail on the glass cup with high probability (0.9). Contrary to when the robot always succeeds, in this case it is actually beneficial for the human to intervene and pick up the glass him- or herself to avoid the large penalty from a likely robot failure.

7.1 Robot Policy with Robot Failures

Figure 4 (bottom row) shows the trust dynamics when the robot fails to accomplish a task. In general, human trust tends to decrease if the robot fails on a task. We recompute the trust-POMDP policy using the trust dynamics model with robot failures. Figure 10 shows the computed trust-POMDP policy and belief updates: the robot starts with the glass cup, since the beginning of the task is when the human is most likely to intervene and not let the robot pick up the glass, which prevents robot failures.

Although this shows that the robot can reason over the human intervention rate to reduce failure, intuitively the robot should also be able to *actively reduce* trust to affect human behavior. There is a range of behaviors that can reduce human trust [43, 44]; however, we focused on active trust reduction through failures. Therefore, we expanded the robot's action space so that it can intentionally fail on any object. Keeping the failure probability for glass at 0.9 and reducing the reward for robot success when picking up the bottles to 0.3 results in the exciting behavior demonstrated in Figure 11.

When following the trust-POMDP policy (Figure 11, top and middle row), the robot attempts to pick up the can first; this is an *information seeking* action that the robot uses to estimate the initial human trust. If the human stays put, the robot infers that the human trust is high, and it will then *fail intentionally* at the bottles to reduce trust, before going for the glass cup. By the time the robot goes for the glass cup, human trust has been reduced sufficiently so that the human is likely to intervene, avoiding failure. However, if the human intervenes, the robot infers that the human trust is already low. The robot then does not need to fail intentionally, as it does not need to reduce human trust any further, so it subsequently goes for the glass cup.

The resulting policy contrasts the policy that the robot follows, if it maximizes human trust instead (Figure 11, bottom row). When following the trust-maximizing policy, the robot starts with the glass. Human trust at the beginning is relatively low, and therefore the human is most

Fig. 11. Sample runs of the performance-maximizing policy (i.e., trust-POMDP, top and middle row) and the trust-maximizing policy (bottom row) when the robot may fail on the glass cup with high probability, and the robot can fail intentionally in any object. The adaptive trust-POMDP policy starts with can object and branches out at $T = 1$. If the human stays put (top row), the robot intentionally fails in the bottles to reduce human trust and maximizes the probability of the human intervening when it goes for the glass at $T = 5$. If the human intervenes (middle row), the robot realizes that the human trust is already low, and it goes for the glass cup at $T = 3$. However, the trust-maximizing policy starts with the glass, since the human is most likely to intervene at the beginning as the human trust is relatively low.

Fig. 12. Top row: Expected trust evolution for all possible human action sequences for the performance-maximizing and trust-maximizing policies. Each action sequence is represented with a line of width proportional to the likelihood of that sequence, based on the learned model. Bottom row: Annotated robot actions and human actions for the top 16 most likely sequences. The robot action is annotated by the object legend (e.g., the red triangle means the robot picks up the can object). The human action is annotated by different line types (e.g., the solid line means the human stays put and the robot succeeds). Overall, the performance-maximizing policy may intentionally fail on the bottle object (bottom left) to calibrate human trust, which results in lower expected human trust than the trust-maximizing policy (top row).

likely to intervene and stop the robot from picking the glass. This prevents significant reduction of human trust, as the robot will fail on the glass with high probability.

7.2 Simulation Results

We further illustrate the difference between the performance-maximizing policy and trust-maximizing policy by simulations. Figure 12 shows the sampled action sequences and trust evolution from the two policies. Overall, the action sequences from the performance-maximizing policy are significantly different from the action sequences from the trust-maximizing policy. The performance-maximizing policy starts with an information gathering action (i.e., pick the can object) and branches at the second step. If the human intervenes, the robot will intentionally fail on the bottles to reduce human trust and increase the probability of human intervention at the glass object (Figure 12, bottom row). This results in lower expected human trust when compared with the trust-maximizing policy (Figure 12, top row).

Figure 13 shows the accumulated rewards as a function of the expected trust for all action sequences from the performance-maximizing policy and the trust-maximizing policy. There are two key observations from this plot. First, maximizing human trust alone can be detrimental to the team performance (i.e., the green circles with large trust value but low accumulated reward). Second, the performance-maximizing policy is able to actively reduce human trust to maximize accumulated reward.

Fig. 13. Scatterplot of mean accumulated reward as a function of human trust over time for all human action sequences of the performance-maximizing policy (blue) and the trust-maximizing policy (green). The radius of each circle is proportional to the likelihood of the corresponding sequence, based on the learned model. The performance-maximizing policy (blue) gradually reduces human trust to maximize the accumulated reward, whereas the trust-maximizing policy (green) focuses on increasing trust.

The mean accumulated reward over 10^4 policy runs for the performance-maximizing policy is -1.33 , compared to -1.67 for the trust-maximizing policy, with a statistically significant difference ($F(1, 19998) = 29.42, p < 0.001$). The effect size $d = 0.1$ indicates a small difference between the two mean accumulated rewards. Overall, the simulation results show that maximizing human trust alone can be suboptimal in the presence of robotic failures.

8 DISCUSSION

Calibrated trust is important in a human-robot team. With the trust-POMDP, the robot can actively modulate human trust to maximize the performance. Our experimental results in a table clearing task show that the trust-POMDP policy calibrates human trust to match it to the robot's manipulation capabilities: if trust is overly low, the robot prioritizes picking up the low-risk objects to increase trust. This results in better performance compared to the myopic robot that ignores trust. However, if trust is overly high, the robot fails intentionally in the low-risk objects. In addition, our results show that maximizing human trust alone can be detrimental to performance in the presence of robotic failures.

Unlike direct communication between the human and the robot, trust builds up over time. The human may still choose to rely on the robot even if the robot indicates its incapability to accomplish a task. For example, although Tesla has told its customers to not overtrust the autopilot system and always put their hands on the wheel (direct communication), the customers may ignore that warning because the autopilot has been performing well for thousands of miles (i.e., trust has been built up).

Computational complexity. Solving a trust-POMDP is essentially solving a standard POMDP model [22], which is computationally expensive in general. The SARSOP algorithm [24, 34] used in this work is the state-of-the-art offline POMDP solver, which computes an optimal action offline for every possible belief state. SARSOP is able to scale up to problems with $O(10^5)$ number of states [34], which roughly equals the table clearing task with nine objects on the table (the number of states in the table clearing task is 4^N , where N is the number of objects).

To scale up trust-POMDP to large robotic tasks, one can apply the online POMDP algorithms. Unlike offline POMDP algorithms, the online POMDP algorithm performs a look-ahead search and computes a best action for the current belief only [38]. After the robot executes the action, the algorithm updates the belief based on the observation received. The process then repeats at the new belief for the next timestep. Online search algorithms scale up by focusing on the current belief only rather than all possible beliefs that the robot may encounter. POMCP [40] and DESPOT [42] are the fastest online POMDP algorithms available today. Both employ the idea of sampling future contingencies. In particular, DESPOT has been successfully implemented for real-time (3hz)

autonomous driving in a crowd [1], where the problem has continuous state and a planning horizon of 90 steps.

Limitations. There are several limitations in our current work. Similar to previous works [9, 46], we modeled trust as a single real-valued latent variable that reflected the capabilities of the entire system. However, a multi-dimensional parameterization of trust that captured the different functions and modes of automation could be a more accurate representation. In addition, the evolution of trust might also depend on the type of motion executed by the robot (e.g., for expressive or deceptive motions [12, 13]). The current trust-POMDP model also assumes static robot capabilities, but a robot's true capabilities may change over time. In fact, the trust-POMDP can be extended to model robot capabilities via additional state variables that affect the state transition dynamics. Furthermore, the reward function is manually specified in this work. However, the reward function may be difficult to specify in practice. One possible way to resolve this is to learn the reward function from human demonstrations (e.g., [33]). Moreover, the work learned a general human behavioral model that applied it to all participants. However, humans might behave differently depending on their trust in the robot. A human behavioral model that adapts to different types of humans is more appropriate. Finally, learning the trust dynamics model relies on direct supervision of trust through a per-action questionnaire, which might become overcomplicated as the problem becomes larger.

9 CONCLUSION

This article presents the trust-POMDP, a computational model for integrating human trust into robot decision making. The trust-POMDP closes the loop between trust models and robot decision making. It enables the robot to infer and influence human trust systematically and to leverage trust for fluid collaboration.

This work points to several directions for further investigation. First, the trust-POMDP is applicable in many real-world robotic tasks. In the current work, we have chosen a simple table clearing task to test the trust-POMDP model, because it allows us to analyze experimentally the core technical issues on the human trust without interference from confounding factors. However, we believe the overall technical approach in our work is general and not restricted to this particular simplified task. What we learned here on the trust-POMDP for a simplified task will be a stepstone toward more complex, large-scale applications. For example, consider a rescue task, where the robot and the human collaborate to rescue survivors in a disaster scene. The robot can quickly search over the disaster scene, but the robot sensors have noises and might fail at certain scenarios. To complete the task successfully, the human and the robot have to collaborate and take advantage of their own strengths. Modeling human trust is important in the rescue task. If the human trust is too low (under-reliance), the rescue will be inefficient as the robot is not well utilized. However, if the human trust is too high (over-reliance), they may miss some survivors as the robot can fail. With the trust-POMDP, the robot can actively modulate human trust to avoid under-reliance and over-reliance.

Second, the trust-POMDP is a generative decision model conditioned explicitly on trust. Explicit trust modeling provides several advantages: it fits better to experimental data (see Section 4.3) and potentially improves the efficiency of learning by reducing sample complexity. Most importantly, the trust model learned on one task may transfer to a related task [41]. This last aspect is another interesting direction for future work.

REFERENCES

- [1] H. Y. Bai, S. J. Cai, D. Hsu, and W. S. Lee. 2015. Intention-aware online POMDP planning for autonomous driving in a crowd. In *Proceedings of the IEEE International Conference on Robotics and Automation*.

- [2] Haoyu Bai, Shaojun Cai, Nan Ye, David Hsu, and Wee Sun Lee. 2015. Intention-aware online POMDP planning for autonomous driving in a crowd. In *Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA'15)*. IEEE, Los Alamitos, CA, 454–460.
- [3] Tirthankar Bandyopadhyay, Kok Sung Won, Emilio Frazzoli, David Hsu, Wee Sun Lee, and Daniela Rus. 2013. Intention-aware motion planning. In *Algorithmic Foundations of Robotics X*. Springer, 475–491.
- [4] Bernard Barber. 1983. *The Logic and Limits of Trust*/Bernard Barber. Rutgers University Press, New Brunswick, NJ.
- [5] Samuel Barrett, Noa Agmon, Noam Hazon, Sarit Kraus, and Peter Stone. 2014. Communicating with unknown teammates. In *Proceedings of the 21st European Conference on Artificial Intelligence*. 45–50.
- [6] Bob Carpenter, Andrew Gelman, Matt Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Michael A. Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2016. Stan: A probabilistic programming language. *Journal of Statistical Software* 20 (2016), 1–37.
- [7] Min Chen, Stefanos Nikolaidis, Harold Soh, David Hsu, and Siddhartha Srinivasa. 2018. Planning with trust for human-robot collaboration. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, New York, NY, 307–315.
- [8] Nathaniel D. Daw, John P. O'Doherty, Peter Dayan, Ben Seymour, and Raymond J. Dolan. 2006. Cortical substrates for exploratory decisions in humans. *Nature* 441, 7095 (2006), 876–879.
- [9] Munjal Desai. 2012. *Modeling Trust to Improve Human-Robot Interaction*. Ph.D. Dissertation. University of Massachusetts Lowell.
- [10] Munjal Desai, Poornima Kaniarasu, Mikhail Medvedev, Aaron Steinfeld, and Holly Yanco. 2013. Impact of robot failures and feedback on real-time trust. In *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction*. IEEE, Los Alamitos, CA, 251–258.
- [11] Munjal Desai, Mikhail Medvedev, Marynel Vázquez, Sean McSheehy, Sofia Gadea-Omelchenko, Christian Bruggeman, Aaron Steinfeld, and Holly Yanco. 2012. Effects of changing reliability on trust of robot systems. In *Proceedings of the 7th Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, New York, NY, 73–80.
- [12] Anca D. Dragan, Rachel M. Holladay, and Siddhartha S. Srinivasa. 2014. An analysis of deceptive robot motion. In *Proceedings of the Robotics: Science and Systems Conference*. 10.
- [13] Anca D. Dragan, Kenton C. T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In *Proceedings of the 2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI'13)*. IEEE, Los Alamitos, CA, 301–308.
- [14] Michael W. Floyd, Michael Drinkwater, and David W. Aha. 2015. Trust-guided behavior adaptation using case-based reasoning. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*. 4261–4267.
- [15] Enric Galceran, Alexander G. Cunningham, Ryan M. Eustice, and Edwin Olson. 2015. Multipolicy decision-making for autonomous driving via changepoint-based behavior prediction. In *Proceedings of the Robotics: Science and Systems Conference*. 2.
- [16] Robert T. Golembiewski and Mark McConkie. 1975. The centrality of interpersonal trust in group processes. *Theories of Group Processes* 131 (1975), 185.
- [17] Robert J. Hall. 1996. Trusting your assistant. In *Proceedings of the 1996 11th Knowledge-Based Software Engineering Conference*. IEEE, Los Alamitos, CA, 42–51.
- [18] Guy Hoffman. 2013. Evaluating fluency in human-robot collaboration. In *Proceedings of the Robotics: Science and Systems Workshop on Human Robot Collaboration*, Vol. 381. 1–8.
- [19] Shervin Javdani, Siddhartha S. Srinivasa, and J. Andrew Bagnell. 2015. Shared autonomy via hindsight optimization. arXiv:1503.07619.
- [20] Jiun-Yin Jian, Ann M. Bisantz, and Colin G. Drury. 2000. Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics* 4, 1 (2000), 53–71.
- [21] Edward E. Jones, Leslie Rock, Kelly G. Shaver, George R. Goethals, and Lawrence M. Ward. 1968. Pattern of performance and ability attribution: An unexpected primacy effect. *Journal of Personality and Social Psychology* 10, 4 (1968), 317.
- [22] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101, 1–2 (1998), 99–134.
- [23] Roderick M. Kramer and Tom R. Tyler. 1995. *Trust in Organizations: Frontiers of Theory and Research*. Sage Publications.
- [24] Hanna Kurniawati, David Hsu, and Wee Sun Lee. 2008. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Proceedings of the Robotics: Science and Systems Conference*.
- [25] John Lee and Neville Moray. 1992. Trust, control strategies and allocation of function in human-machine systems. *Ergonomics* 35, 10 (1992), 1243–1270.
- [26] John D. Lee and Katrina A. See. 2004. Trust in automation: Designing for appropriate reliance. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 46, 1 (2004), 50–80.

- [27] Owen Macindoe, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. 2012. POMCoP: Belief space planning for sidekicks in cooperative games. In *Proceedings of the 8th Artificial Intelligence and Interactive Digital Entertainment Conference*.
- [28] John E. Mathieu, Tonia S. Heffner, Gerald F. Goodwin, Eduardo Salas, and Janis A. Cannon-Bowers. 2000. The influence of shared mental models on team process and performance. *Journal of Applied Psychology* 85, 2 (2000), 273.
- [29] Roger C. Mayer, James H. Davis, and F. David Schoorman. 1995. An integrative model of organizational trust. *Academy of Management Review* 20, 3 (1995), 709–734.
- [30] Bonnie Marlene Muir. 1990. *Operators' Trust in and Use of Automatic Controllers in a Supervisory Process Control Task*. University of Toronto.
- [31] Stefanos Nikolaidis, David Hsu, and Siddhartha Srinivasa. 2017. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *International Journal of Robotics Research* 36, 5–7 (2017), 618–634.
- [32] Stefanos Nikolaidis, Anton Kuznetsov, David Hsu, and Siddhartha Srinivasa. 2016. Formalizing human-robot mutual adaptation: A bounded memory model. In *Proceedings of the 2016 11th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI'16)*. IEEE, Los Alamitos, CA, 75–82.
- [33] Stefanos Nikolaidis, Ramya Ramakrishnan, Keren Gu, and Julie Shah. 2015. Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In *Proceedings of the 2015 10th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI'15)*. ACM, New York, NY, 189–196.
- [34] Sylvie C. W. Ong, Shao Wei Png, David Hsu, and Wee Sun Lee. 2010. Planning under uncertainty for robotic tasks with mixed observability. *International Journal of Robotics Research* 29, 8 (2010), 1053–1068.
- [35] Alyssa Pierson and Mac Schwager. 2016. Adaptive inter-robot trust for robust multi-robot sensor coverage. In *Robotics Research*. Springer, 167–183.
- [36] Charles Pippin and Henrik Christensen. 2014. Trust modeling in multi-robot patrolling. In *Proceedings of the 2014 IEEE International Conference on Robotics and Automation (ICRA '14)*. IEEE, Los Alamitos, CA, 59–66.
- [37] John K. Rempel, John G. Holmes, and Mark P. Zanna. 1985. Trust in close relationships. *Journal of Personality and Social Psychology* 49, 1 (1985), 95.
- [38] Stephane Ross, Joelle Pineau, Sebastien Paquet, and Brahim Chaib-Draa. 2008. Online planning algorithms for POMDPs. *Journal of Artificial Intelligence Research* 32 (2008), 663–704.
- [39] Dorsa Sadigh, Shankar Sastry, Sanjit A. Seshia, and Anca D. Dragan. 2016. Planning for autonomous cars that leverages effects on human actions. In *Proceedings of the Robotics: Science and Systems Conference (RSS'16)*.
- [40] D. Silver and J. Veness. 2010. Monte-Carlo planning in large POMDPs. In *Advances in Neural Information Processing Systems*.
- [41] Harold Soh, Pan Shu, Min Chen, and David Hsu. 2018. The transfer of human trust in robot capabilities across tasks. arXiv:1807.01866.
- [42] A. Somani, N. Ye, D. Hsu, and W. S. Lee. 2013. DESPOT: Online POMDP planning with regularization. In *Advances in Neural Information Processing Systems*.
- [43] Rik van den Brule, Ron Dotsch, Gijsbert Bijlstra, Daniel H. J. Wigboldus, and Pim Haselager. 2014. Do robot performance and behavioral style affect human trust? *International Journal of Social Robotics* 6, 4 (2014), 519–531.
- [44] Ning Wang, David V. Pynadath, and Susan G. Hill. 2016. Trust calibration within a human-robot team: Comparing automatically generated explanations. In *Proceedings of the 2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI'16)*. IEEE, Los Alamitos, CA, 109–116.
- [45] Yaqi Xie, Indu Prasad, Desmond Ong, David Hsu, and Harold Soh. 2019. Robot capability and intention in trust-based decisions across tasks. In *Proceedings of the 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI'19)*. ACM, New York, NY.
- [46] Anqi Xu and Gregory Dudek. 2015. Optimo: Online probabilistic trust inference model for asymmetric human-robot collaborations. In *Proceedings of the 2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI'15)*. ACM, New York, NY, 221–228.
- [47] Anqi Xu and Gregory Dudek. 2016. Towards modeling real-time trust in asymmetric human-robot collaborations. In *Robotics Research*. Springer, 113–129.
- [48] Jessie Yang, Vaibhav Unhelkar, Kevin Li, and Julie Shah. 2017. Evaluating effects of user experience and system transparency on trust in automation. In *Proceedings of the 2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI'17)*.

Received August 2018; revised May 2019; accepted August 2019