

Enabling Team of Teams: A Trust Inference and Propagation (TIP) Model in Multi-Human Multi-Robot Teams

Yaohui Guo, X. Jessie Yang, and Cong Shi
 Industrial and Operations Engineering
 University of Michigan
 Ann Arbor, Michigan, USA
 {yaohuig, xijyang, shicong}@umich.edu

Abstract—Trust has been identified as a central factor for effective human-robot teaming. Existing literature on trust modeling predominantly focuses on dyadic human-autonomy teams where one human agent interacts with one robot. There is little, if not no, research on trust modeling in teams consisting of multiple human agents and multiple robotic agents. To fill this research gap, we present the trust inference and propagation (TIP) model for trust modeling in multi-human multi-robot teams. In a multi-human multi-robot team, we postulate that there exist two types of experiences that a human agent has with a robot: direct and indirect experiences. The TIP model presents a novel mathematical framework that explicitly accounts for both types of experiences. To evaluate the model, we conducted a human-subject experiment with 15 pairs of participants ($N = 30$). Each pair performed a search and detection task with two drones. Results show that our TIP model successfully captured the underlying trust dynamics and significantly outperformed a baseline model. To the best of our knowledge, the TIP model is the first mathematical framework for computational trust modeling in multi-human multi-robot teams.

I. INTRODUCTION

A human agent’s trust in an autonomous/robotic agent, defined as “the attitude that an agent will help achieve an individual’s goals in situations characterized by uncertainty and vulnerability [1]”, is a central factor for effective human-robot interaction (HRI) [2]–[4]. Optimal interaction can be achieved only when an appropriate level of trust is established between the human and the robotic agents [5], [6]. Despite the extensive research efforts over the past thirty years, existing literature is predominantly focused on trust modeling in dyadic human-robot teams where one human agent interacts with one robot [7]. There is little, if not no, research on trust modeling in teams consisting of multiple human agents and multiple robots.

Consider a scenario where two human agents (figure 1), x and y , and two robots, A and B , are to perform a task. The four agents are allowed to form sub-teams to enhance task performance (e.g., maximizing throughput and minimizing task completion time). For instance, they could initially form two dyadic human-robot teams to complete the first part of the task, merge to complete the second part and split again with a different configuration to complete the third part of the

task, and so on. This scenario illustrates a new organizational model known as “team of teams [8], [9]” in which the team composition is fluid and team members come and go as the nature of the problem changes.

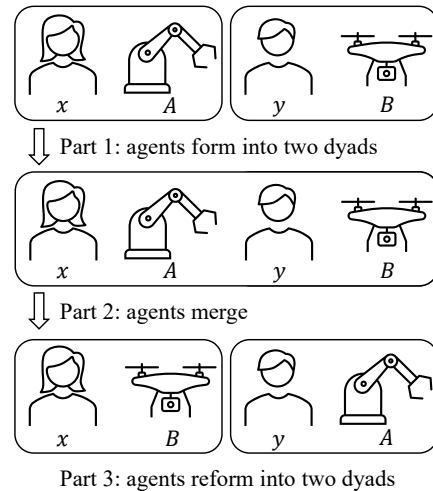


Fig. 1: Four agents can form sub-teams. In part 1, human x and robot A form a dyad, and human y and robot B form a dyad. In part 2, two dyads merge. In part 3, human x and robot B form a dyad, and human y and robot A form a dyad.

In this scenario, we postulate that there exist two types of experiences that a human agent has with a robot: *direct and indirect experiences*. Direct experience, by its name, means that a human agent’s interaction with a robot is by him-/her-self; indirect experience means that a human agent’s interaction with a robot is mediated by another party. Considering the third part of the task (see figures 1 and 2), human x works directly with robot B (i.e., direct experience). Even though there is no direct interaction between x and A in part 3, we postulate that x could still update his or her trust in A by learning her human teammate y ’s experience with A , i.e., y ’s direct experience with A becomes x ’s *indirect* experience with A , based on which x can update her trust in A , $t^{x,A}$. Essentially, y ’s trust in A *propogates* to x .




Fig. 2: An arrow points from a trustor to a trustee, representing the trust $t^{\text{trustor}, \text{trustee}}$. Human x updates her trust in robot B based on direct experience. Even though x does not have direct interaction with A , x could still update her trust toward A through a third party, y .

Under the direct and indirect experience framework, prior work on trust modeling in dyadic human-robot teams can be regarded as examining how *direct* experience influences a person’s trust in a robot. In multi-human multi-robot teams, we postulate that *both direct and indirect experiences drive a human agent’s trust in a robot*.

In this study, we develop the **Trust Inference and Propagation (TIP)** model for multi-human multi-robot teams. The proposed model explicitly accounts for the direct and indirect experiences a human agent may have with a robot. We examine trust dynamics under the TIP framework and prove theoretically that trust converges after repeated (direct and indirect) interactions. To evaluate the proposed TIP model, we conducted a human-subject experiment with 15 pairs of participants ($N = 30$). Each pair worked with two drones to perform a threat detection task for 15 sessions. We compared the TIP model (i.e., accounts for both the direct and indirect experiences) and a direct-experience-only model (i.e., only accounts for the direct experience a human agent has with a robot). Results show that the TIP model successfully captures people’s trust dynamics with a significantly smaller root-mean-square error (RMSE) compared to the direct-experience-only model.

The key contribution of this work is three-fold:

- To the best of our knowledge, the proposed TIP model is the first mathematical framework for computational trust modeling in multi-human multi-robot teams. The TIP model accounts for both the direct and indirect experiences (through trust propagation) a human agent has with a robot in multi-human multi-robot teams. As a result, the TIP model is well-suited for trust estimation in networks involving multiple humans and robots.
- We prove theoretically that trust converges to the unique equilibrium in probability after repeated direct and indirect interactions under our TIP framework. Such an equilibrium can also be efficiently computed.
- We conduct a human-subject experiment to assess the TIP model. Results reveal the superior performance of the TIP model in capturing trust dynamics in a multi-human multi-robot team.

This paper is organized as follows. Section II presents related work, including trust modeling in dyadic human-robot teams and reputation/trust management in e-commerce. In Section III, we describe the mathematical formulation of the TIP model and examine its behavior under different types of interactions.

Section IV presents the human-subject study. In Section V, we present and discuss the results. Section VI concludes the paper.

II. RELATED WORK

In this section, we review two bodies of research motivating the present study: the extensive literature on trust in dyadic human-robot teams and the literature on reputation/trust management. The latter is a research topic in computer science that shares commonalities with the underlying research question of trust modeling in multi-human multi-robot teams.

A. Trust Modeling in Dyadic Human-robot Interaction

Trust in autonomous/robotic agents attracts research attention from multiple disciplines. One line of research is to identify factors influencing a human’s trust in autonomy/robots and quantify their effects. These factors can be categorized into human-related factors such as personality [10], robot-related factors such as reliability [11], [12] and transparency [13], [14], and task-related factors such as task emergency [15]. For a review of the factors, see [6]. More recently, another line of research has emerged that focuses on understanding the dynamics of trust formation and evolution when a person interacts with autonomy repeatedly [3], [16], [17]. Empirical studies have investigated how trust strengthens or decays due to moment-to-moment interactions with autonomy [3], [18]–[20]. Based on the empirical research, three major properties of trust dynamics have been identified and summarized, namely *continuity*, *negativity bias*, and *stabilization* [17], [21]

Acknowledging that trust is a dynamic variable, several computational trust models in dyadic human-robot teams have been developed [17], [22]–[24]. Notably, Xu and Dudek [23] proposed the online probabilistic trust inference model (OPTIMo) utilizing Bayesian networks to estimate human trust based on the autonomous agent’s performance and human behavioral signals. Guo and Yang [17] modeled trust as a Beta random variable parameterized by positive and negative interaction experience a human agent has with a robotic agent. Soh et al. [25] proposed a Bayesian model which combines Gaussian processes and recurrent neural networks to predict trust over different tasks. For a detailed review, refer to [26].

B. Reputation/Credential Management

Despite limited research on trust modeling in multi-human multi-robot teams, insights can be drawn from studies on reputation management. In consumer-to-consumer electronic marketplaces like eBay, reputation systems play a crucial role in generating trust among buyers to facilitate transactions with unknown sellers [27]. These systems can be categorized as centralized, where reputation values are stored centrally, representing the overall trustworthiness of sellers, or decentralized, where buyers maintain their evaluation scores privately [28]. In decentralized systems, a propagation mechanism allows buyers to obtain reputation values, even in the absence of prior transactions. Various propagation mechanisms have been developed, such as subjective logic integrated into the Beta reputation management system [29], [30] or the concept

of "witness reputation" in the FIRE reputation management model [31], facilitating the transfer of reputation scores among agents in a network. These propagation mechanisms provide valuable insights into modeling trust update through indirect experience in HRI. Yet, their direct application to HRI settings is impeded due to the distinct characteristics of human trust in robots, as reviewed in Section II-A.

III. MATHEMATICAL MODEL

We present the TIP model in this section. Our key motivation is to develop a fully computational trust inference and propagation model that works in general multi-human multi-robot settings. First, we discuss the assumptions and introduce the mathematical formulation. Second, we examine the behavior of the model under repeated human-robot interactions. Finally, we present the parameter inference method and trust estimation using the TIP model.

A. Assumptions

We make three major assumptions in the context of HRI. First, we assume that *each human agent communicates trust as a single-dimensional value*. In some prior work, trust is represented as a tuple. For example, in [30], trust is represented as a triplet, i.e., belief, disbelief, and uncertainty. Although a multi-dimensional representation conveys more information, our study as well as some prior studies show that a one-dimensional representation of trust suffices in capturing trust evolution [17], [22]–[24], [32]. Moreover, querying a single-dimension trust value increases operational feasibility because keeping track of multiple numbers adds unnecessary cognitive load and may not be pragmatic for non-experts. Therefore, we assume a simple one-dimension form of trust in this study.

Second, we assume that *the human agents are cooperative*, i.e., they are honest and willing to share their trust in a robot truthfully with their human teammates.

Third, we take an ability/performance-centric view of trust and assume that a human agent's trust in a robot is primarily driven by the ability or performance of the robot. This ability/performance-centric view has been widely used in prior research for modeling trust in task-oriented HRI contexts (i.e., a robot is to perform a specific task) [17], [22], [23].

We discuss the limitations of the assumptions in Section VI.

B. Proposed Model

Trust as a Beta random variable. We take a probabilistic view to model trust as in [17]. At time k , the trust $t_k^{a,b}$ that a human agent a feels toward another agent b follows a Beta distribution, i.e.,

$$t_k^{a,b} \sim \text{Beta}(\alpha_k^{a,b}, \beta_k^{a,b}), \quad (1)$$

where $\alpha_k^{a,b}$ and $\beta_k^{a,b}$ are the positive and negative experiences a had about b up to time k , respectively, $k = 0, 1, 2, \dots$. When $k = 0$, $\alpha_0^{a,b}$ and $\beta_0^{a,b}$ represent the prior experiences that a has before any interaction with b . The expected trust is given by

$$\mu_k^{a,b} = \alpha_k^{a,b} / (\alpha_k^{a,b} + \beta_k^{a,b}). \quad (2)$$

Here we note that $t_k^{a,b}$ is the queried trust given by the agent a , which has some randomness due to subjectivity, while $\mu_k^{a,b}$ is the expected trust determined by the experiences.

Trust update through direct experience. Similar to [17], we update the experiences through direct interaction at time k by setting

$$\begin{aligned} \alpha_k^{a,b} &= \alpha_{k-1}^{a,b} + s^{a,b} \cdot p_k^b \\ \beta_k^{a,b} &= \beta_{k-1}^{a,b} + f^{a,b} \cdot \bar{p}_k^b \end{aligned} \quad (3)$$

Here p_k^b and \bar{p}_k^b are the measurements of b 's success and failure during time k , respectively; $s^{a,b}$ and $f^{a,b}$ are a 's unit experience gains with respect to success or failure of b . We require $s^{a,b}$ and $f^{a,b}$ to be positive to ensure that cumulative experiences are non-decreasing. The updated trust $t_k^{a,b}$ follows the distribution $\text{Beta}(\alpha_k^{a,b}, \beta_k^{a,b})$.

Trust update through indirect experience propagation. Let x and y denote two human agents and let A denote a robot agent, as illustrated in figure 2. At time k , y communicates his or her trust $t_k^{y,A}$ in A with x , and then x updates his or her experiences through indirect interaction by

$$\begin{aligned} \alpha_k^{x,A} &= \alpha_{k-1}^{x,A} + \hat{s}^{x,A} \cdot t_k^{x,y} \cdot \left[t_k^{y,A} - t_{k-1}^{x,A} \right]^+ \\ \beta_k^{x,A} &= \beta_{k-1}^{x,A} + \hat{f}^{x,A} \cdot t_k^{x,y} \cdot \left[t_{k-1}^{x,A} - t_k^{y,A} \right]^+ \end{aligned} \quad (4)$$

where the superscript '+' means taking the positive part of the corresponding number, i.e., $t^+ = \max\{0, t\}$ for a real number t , and $t_k^{x,A} \sim \text{Beta}(\alpha_k^{x,A}, \beta_k^{x,A})$.

The intuition behind this model is that x needs to reason upon $t_k^{y,A}$, i.e., y 's trust toward A . First, x compares y 's trust $t_k^{y,A}$ with his or her previous trust $t_{k-1}^{x,A}$. Let $\Delta t := t_k^{y,A} - t_{k-1}^{x,A}$ be the difference. If $\Delta t \geq 0$, x gains positive indirect experience about A , which amounts to the product of the trust difference Δt , a coefficient $\hat{s}^{x,A}$, and a discounting factor $t_k^{x,y}$, i.e., x 's trust in y ; if $\Delta t < 0$, then x gains negative indirect experience about A , which is defined similarly.

C. Asymptotic Behavior under Repeated Interactions

We examine the behavior of the proposed model under both direct and indirect trust updates. Consider a scenario where human agents x and y take turns working with robot A repetitively. Suppose each x 's turn contains m interactions while each y 's turn contains n interactions; and, after each interaction, the agent who works directly with A informs the other agent of his or her trust in A . Figure 3 illustrates the interaction process. In addition, we assume that robot A has constant reliability r , i.e., A 's performance measure are $p_k^A = r$ and $\bar{p}_k^A = \bar{r}$, for $k = 1, 2, \dots, K$, where $\bar{r} := 1 - r$, and x has constant trust $t^{x,y}$ in y . To avoid triviality, we exclude the case when $m = n = 0$ (where no interactions occur). Without loss of generality, we assume $m > 0$ and $n \geq 0$. (The case $m \geq 0$ and $n > 0$ is symmetric.)

We have the following main result on the asymptotic behavior of $t_k^{x,A}$ and $t_k^{y,A}$.




Fig. 3: x and y take turns to interact with A .

Theorem 1. When $m > 0$ and $n \geq 0$, $t_k^{x,A}$ and $t_k^{y,A}$ converge in probability (i.p.) respectively, i.e., there exists t^x and t^y such that, for any $\epsilon > 0$,

$$\lim_{k \rightarrow \infty} \Pr \left(\left| t_k^{x,A} - t^x \right| > \epsilon \right) = 0$$

and

$$\lim_{k \rightarrow \infty} \Pr \left(\left| t_k^{y,A} - t^y \right| > \epsilon \right) = 0.$$

Theorem 1 exhibits that, under alternating interactions with the robot, both agents' trust will stabilize and converge after sufficiently many interactions. The next result gives an exact method to compute the limiting equilibrium.

Theorem 2. The equilibrium t^x and t^y in Theorem 1 satisfy

$$S^x \frac{1 - t^x}{t^x} = \hat{F}^x (t^x - t^y) + F^x \quad \text{and}$$

$$F^y \frac{t^y}{1 - t^y} = \hat{S}^y (t^x - t^y) + S^y, \quad (5)$$

if $S^x F^y \geq F^x S^y$; otherwise, they satisfy

$$F^x \frac{t^x}{1 - t^x} = \hat{S}^x (t^y - t^x) + S^x \quad \text{and}$$

$$S^y \frac{1 - t^y}{t^y} = \hat{F}^y (t^y - t^x) + F^y, \quad (6)$$

where $\hat{S}^x = nt^{x,y} \hat{s}^{x,A}$, $\hat{F}^x = nt^{x,y} \hat{f}^{x,A}$, $S^x = ms^{x,A}r$, $F^x = mf^{x,A}\bar{r}$, $\hat{S}^y = mt^{y,x} \hat{s}^{y,A}$, $\hat{F}^y = mt^{y,x} \hat{f}^{y,A}$, $S^y = ns^{y,A}r$, and $F^y = nf^{y,A}\bar{r}$.

The capitalized variables in Theorem 2 are related to the average experience gains in the long run, e.g., S^x is x 's direct positive experience gain after each m direct update. The condition $S^x F^y \geq F^x S^y$ can be interpreted as follows: compared with y , x tends to have a higher trust gain in A after each turn via direct experience. Note that t^x and t^y can be computed exactly by solving a cubic equation or readily approximated by Newton's method. Details are given in the appendix.

A special case is when $n = 0$, i.e., agent x only updates trust in A via direct experience, and agent y only updates trust via indirect experience. Theorem 2 leads to the following corollary with a closed-form equilibrium:

Corollary 1. When $m > 0$ and $n = 0$, x 's trust $t_k^{x,A}$ in A converges to $t^x = \frac{s^{x,A}r}{f^{x,A}\bar{r} + s^{x,A}r}$ in probability, i.e., for any $\epsilon > 0$,

$$\lim_{k \rightarrow \infty} \Pr \left(\left| t_k^{x,A} - \frac{s^{x,A}r}{f^{x,A}\bar{r} + s^{x,A}r} \right| > \epsilon \right) = 0.$$

The difference between $t_k^{x,A}$ and $t_k^{y,A}$ converge to 0 in probability, i.e., for any $\epsilon > 0$,

$$\lim_{k \rightarrow \infty} \Pr \left(\left| t_k^{x,A} - t_k^{y,A} \right| > \epsilon \right) = 0.$$

Equivalently, we have $t^x = t^y$ in Theorem 2.

Corollary 1 implies that under direct-only trust updates, x 's trust will stabilize around the following closed-form

$$\frac{s^{x,A}r}{f^{x,A}\bar{r} + s^{x,A}r},$$

which is determined by x 's unit experience gains $s^{x,A}$ and $f^{x,A}$ via direct trust, and the robot's reliability r ; moreover, under indirect-only updates, y 's trust will converge to x 's trust.

D. Parameter Inference

The proposed model characterizes a human agent's trust in a robot by six parameters. For instance, the parameter of x on robot A , which is defined as

$$\theta^{x,A} = \left(\alpha_0^{x,A}, \beta_0^{x,A}, s^{x,A}, f^{x,A}, \hat{s}^{x,A}, \hat{f}^{x,A} \right), \quad (7)$$

includes x 's prior experiences $\alpha_0^{x,A}$ and $\beta_0^{x,A}$, the unit direct experience gains $s_A^{x,A}$ and $f_A^{x,A}$, and the unit indirect experience gains $\hat{s}_A^{x,A}$ and $\hat{f}_A^{x,A}$. Denote the indices of x 's direct and indirect interactions with A up to time k as D_k and \bar{D}_k . We can compute $\alpha_k^{x,A}$ and $\beta_k^{x,A}$, according to Eqs. (3) and (4), as

$$\alpha_k^{x,A} = \alpha_0^{x,A} + s^{x,A} \sum_{j \in D_k} p_j^A + \hat{s} \sum_{j \in \bar{D}_k} t_j^{x,y} \left[t_j^{x,A} - t_{j-1}^{x,A} \right]^+$$

$$\beta_k^{x,A} = \beta_0^{x,A} + f^{x,A} \sum_{j \in D_k} \bar{p}_j^A + \hat{f} \sum_{j \in \bar{D}_k} t_j^{x,y} \left[t_{j-1}^{x,A} - t_j^{y,A} \right]^+ \quad (8)$$

We compute the optimal parameter $\theta_*^{x,A}$ by maximum likelihood estimation (MLE), i.e.,

$$\theta_*^{x,A} = \arg \max \log \Pr \left(\text{data} \mid \theta^{x,A} \right)$$

$$= \arg \max \sum_{k=0}^K \log \text{Beta} \left(t_k^{x,A} \mid \alpha_k^{x,A}, \beta_k^{x,A} \right).$$

Specifically, the problem of estimating x 's parameter $\theta_*^{x,A}$ on robot A is formulated as follows: given x 's full trust history in A , $\{t_k^{x,A}\}_{k=0,1,\dots,K}$, A 's performance history during x 's direct trust update in A , $\{(p_k^A, \bar{p}_k^A)\}_{k \in D_K}$, x 's trust in y during x 's indirect trust update in A , $\{t_k^{x,y}\}_{k \in \bar{D}_K}$, and y 's trust in A during x 's indirect trust update in A , $\{t_k^{y,A}\}_{k \in \bar{D}_K}$, we compute the parameter $\theta_*^{x,A}$ that maximizes the log likelihood function

$$H(\theta^{x,A}) := \sum_{k=0}^K \log \text{Beta} \left(t_k^{x,A} \mid \alpha_k^{x,A}, \beta_k^{x,A} \right), \quad (9)$$

where $\alpha_k^{x,A}$ and $\beta_k^{x,A}$ are defined in Eq. (8).

We note that $\log \text{Beta}(t_k^{x,A} \mid \alpha_k^{x,A}, \beta_k^{x,A})$ is concave in $\theta^{x,A}$ because it is concave in $(\alpha_k^{x,A}, \beta_k^{x,A})$ and $\alpha_k^{x,A}$ and $\beta_k^{x,A}$ are non-decreasing linear functions of $\theta^{x,A}$. Consequently, $H(\theta^{x,A})$ is concave in $\theta^{x,A}$ since it is a summation of several concave functions. Therefore, we can run the gradient descent method to compute the optimal parameters.

Now we explicitly give the formulas for the gradient descent method. By expressing the probability density function of Beta

random variables in terms of Gamma functions, we can rewrite Eq. (9) as

$$\begin{aligned} & H(\boldsymbol{\theta}^{x,A}) \\ &= \sum_{k=0}^K \left[\log \Gamma(\alpha_k^{x,A} + \beta_k^{x,A}) - \log \Gamma(\alpha_k^{x,A}) - \log \Gamma(\beta_k^{x,A}) \right. \\ & \quad \left. + (\alpha_k^{x,A} - 1) \log t_k^{x,A} + (\beta_k^{x,A} - 1) \log(1 - t_k^{x,A}) \right], \end{aligned}$$

where $\Gamma(\cdot)$ stands for the Gamma function. Define the following variables:

$$\begin{aligned} P_k &:= \sum_{j \in D_k} p_j^A, \quad Q_k := \sum_{j \in \bar{D}_k} t_j^{x,y} \left[t_j^{y,A} - t_{j-1}^{x,A} \right]^+, \\ \bar{P}_k &:= \sum_{j \in D_k} \bar{p}_j^A, \quad \bar{Q}_k := \sum_{j \in \bar{D}_k} t_j^{x,y} \left[t_{j-1}^{x,A} - t_j^{y,A} \right]^+. \end{aligned}$$

Then (8) becomes

$$\begin{aligned} \alpha_k^{x,A} &= \alpha_0^{x,A} + s^{x,A} P_k + \hat{s}^{x,A} Q_k \\ \beta_k^{x,A} &= \beta_0^{x,A} + f^{x,A} \bar{P}_k + \hat{f}^{x,A} \bar{Q}_k. \end{aligned}$$

Calculation shows the gradient can be written as

$$\nabla H(\boldsymbol{\theta}^{x,A}) = \sum_{k=0}^K \mathbf{C}_k \mathbf{v}_k, \quad (10)$$

where

$$\mathbf{C}_k = \begin{bmatrix} 1 & -1 & 0 & 1 & 0 \\ 1 & 0 & -1 & 0 & 1 \\ P_k & -P_k & 0 & P_k & 0 \\ \bar{P}_k & 0 & -\bar{P}_k & 0 & \bar{P}_k \\ Q_k & -Q_k & 0 & Q_k & 0 \\ \bar{Q}_k & 0 & -\bar{Q}_k & 0 & \bar{Q}_k \end{bmatrix} \quad (11)$$

and

$$\mathbf{v}_k = \begin{bmatrix} \psi \left(\alpha_k^{x,A} + \beta_k^{x,A} \right) \\ \psi \left(\alpha_k^{x,A} \right) \\ \psi \left(\beta_k^{x,A} \right) \\ \log t_k^{x,A} \\ \log \left(1 - t_k^{x,A} \right) \end{bmatrix}. \quad (12)$$

Here ψ is the digamma function. Note that \mathbf{C}_k is constant throughout the gradient descent while \mathbf{v}_k needs to be computed in every iteration.

E. Trust Estimation

In real HRI scenarios, querying human trust after every interaction is impractical as it introduces extra workload and reduces collaboration efficiency. Instead, we consider the case when human trust is only queried after some, but not all, of the interactions. In particular, we are interested in referring the model parameter $\boldsymbol{\theta}^{x,A}$ defined in Eq. (7) with missing trust values and estimating these missing values with the TIP model.

Specifically, the input of the trust estimation problem is the same as the parameter inference problem in Section III-D, except that $t_u^{x,A}$, $t_u^{x,y}$, and $t_u^{y,A}$ are missing for $u \in U$, where U

is the collection of interactions without trust ratings. We assume $0 \notin U$, that is, the initial trust ratings, $t_0^{x,y}$, $t_0^{y,A}$, and $t_0^{x,A}$, are known. The optimal parameter is defined as the maximizer of the log-likelihood given the available data:

$$H_U(\boldsymbol{\theta}^{x,A}) := \sum_{k \in \{0, \dots, K\} \setminus U} \log \text{Beta} \left(t_k^{x,A} \middle| \alpha_k^{x,A}, \beta_k^{x,A} \right).$$

Eq. (8) implies that computing the experiences $\alpha_k^{x,A}$ and $\beta_k^{x,A}$ relies on the trust ratings $t_j^{x,y}$, $t_j^{y,A}$, and $t_j^{x,A}$. We approximate them by the following recursive relations:

$$\hat{t}_j^{x,y} = t_j^{x,y}, \quad \hat{t}_j^{y,A} = t_j^{y,A}, \quad \text{and} \quad \hat{t}_j^{x,A} = t_j^{x,A},$$

for $j \notin U$;

$$\hat{t}_j^{x,y} = t_{j'}^{x,y}, \quad \hat{t}_j^{y,A} = t_{j'}^{y,A}, \quad \text{and} \quad \hat{t}_j^{x,A} = t_{j'}^{x,A}, \quad (13)$$

for $j \in U$, where $j' = \max\{0, 1, \dots, j-1\} \setminus U$. In other words, we use the trust rating from the most recent interactions to approximate the missing values. We note that the index j' is well defined in Eq. (13) since we assume the initial trust ratings are known. Now, we can compute $\alpha_k^{x,A}$ and $\beta_k^{x,A}$ by the approximated trust values as follows

$$\begin{aligned} \alpha_k^{x,A} &= \alpha_0^{x,A} + s^{x,A} \sum_{j \in D_k} p_j^A + \hat{s} \sum_{j \in \bar{D}_k} \hat{t}_j^{x,y} \left[\hat{t}_j^{y,A} - \hat{t}_{j-1}^{x,A} \right]^+ \\ \beta_k^{x,A} &= \beta_0^{x,A} + f^{x,A} \sum_{j \in D_k} \bar{p}_j^A + \hat{f} \sum_{j \in \bar{D}_k} \hat{t}_j^{x,y} \left[\hat{t}_{j-1}^{x,A} - \hat{t}_j^{y,A} \right]^+ \end{aligned} \quad (14)$$

Similar to maximizing H , we can apply the gradient descent method to find the maximizer $\boldsymbol{\theta}_*^{x,A}$ of H_U . The gradient ∇H_U can be computed in the same way as Eq. (10) except that the summation is over $\{0, \dots, K\} \setminus U$ instead of $\{0, \dots, K\}$, i.e.,

$$\nabla H_U(\boldsymbol{\theta}^{x,A}) = \sum_{k \in \{0, \dots, K\} \setminus U} \mathbf{C}_k \mathbf{v}_k,$$

where \mathbf{C}_k and \mathbf{v}_k are defined in Eqs. (11) and (12) and computed with the estimated trust values.

By substituting $\boldsymbol{\theta}_*^{x,A}$ to Eq. (14), we can approximate the experiences and further estimate the missing trust rating $t_u^{x,A}$ by the expectation $\mu_u^{x,A} = \frac{\alpha_u^{x,A}}{\alpha_u^{x,A} + \beta_u^{x,A}}$ for $u \in U$.

IV. HUMAN SUBJECT STUDY

We conducted a human-subject experiment with 30 participants to evaluate the TIP model. The experiment, inspired by [20], simulated a search and detection task where two human agents work with two smart drones to search for threats at multiple sites.

A. Participants

A total of $N = 30$ participants (average age = 25.3 years, SD = 4.3 years, 16 females, 14 males) with normal or corrected-to-normal vision formed 15 teams and participated in the experiment. Each participant received a base payment of \$15 and a bonus of up to \$10 depending on their team performance.

B. Experimental Task and Design


In the experiment, a pair of participants performed a simulated threat detection task with two assistant drones for $K = 15$ sessions on two separate desktop computers. At each session, each participant was assigned one drone and worked on the detection tasks. After the session, they were asked to report their trust in each drone and their trust in their human teammate. For clarity, we named the two drones A and B and colored them in red and blue, respectively; and we denoted the participants as x and y . A trust rating is denoted as $t_k^{a,b}$, where the superscript $a \in \{x, y\}$ stands for the trustor, the superscript $b \in \{x, y, A, B\}$ stands for the trustee, and the subscript k is the session index. For example, $t_2^{x,A}$ is person x 's trust in drone A after the 2nd session. The range of a trust rating is $[0, 1]$, where 0 stands for “(do) not trust at all” and 1 stands for “trust completely”. The flow of the experimental task is illustrated in figure 4a.

Initial trust rating. At the start, each participant gave their initial trust in the two drones based on their prior experience with automation/robots. Additionally, they gave their initial trust in each other. These trust ratings were indexed by 0, e.g., x 's initial trust rating in A was denoted as $t_0^{x,A}$.

Robot assignment. At each session, each participant was randomly assigned one drone as his or her assistant robot, as shown in figure 5.

Search and detection task. Each session consisted of 10 locations to detect. As shown in figure 4b, four views were present at each location. If a threat, which appeared like a combatant, was in any of the views, the participant should click the ‘Danger’ button; otherwise, they should click the ‘Clear’ button. Meanwhile, his or her drone would assist by highlighting a view if the drone detected a threat there. In addition, a 3-second timer was set for each location. If a participant did not click either button before the timer counted down to zero, the testbed would move to the next location automatically. After all the 10 locations, an end-of-session screen was shown, displaying how many correct choices the participant and the drone had made in the current session. Correct choices mean correctly identifying threats or declaring ‘Clear’ within 3 seconds.

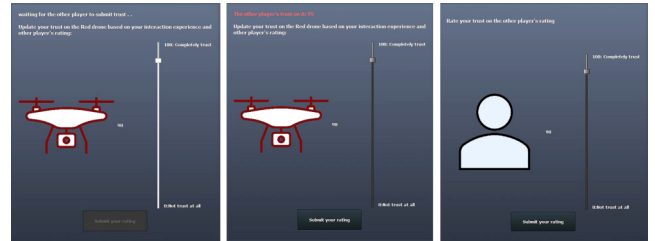
Trust rating. After each session, each participant reported three trust values. First, each participant updated his or her trust in the drone s/he just worked with, i.e., through direct experience, based on the drone’s detection ability. Next, through a server (see figure 5), each participant communicated their trust in the drone s/he just worked with to their human teammate. After that, each participant updated his or her trust in the other player’s drone (i.e., through indirect experience). Note that only trust ratings were communicated and drones’ performances were not. Finally, each participant updated his or her trust in the human teammate based on the teammate’s ability to rate trust in the drones accurately. Hence, after the k th session, there would be 6 additional self-reported trust values, $t_k^{x,A}$, $t_k^{x,B}$, $t_k^{y,A}$, $t_k^{y,B}$, $t_k^{x,y}$, and $t_k^{y,x}$. An illustration of the rating



(a) Experiment process.



(b) Task interface. The drone will highlight the potential threat in bright red. The participant is asked to click the ‘Danger’ button if a threat is present and to click the ‘Clear’ button otherwise.



(c) Trust rating interface. Suppose player x was assigned drone B (blue drone). The system first prompted x to rate his or her trust in B based on direct experience and then to update trust in A via indirect experience. The interface of rating A first showed the left figure, where the prompt read “waiting for the other player to submit trust”. After the other player y rated his or her trust in A , the prompt changed the text to “The other player’s trust on A : 95”, and asked for player x ’s rating: “Update your trust in the red drone based on your interaction experience and the other player’s rating”, as shown in the middle figure. Afterward, x would rate trust in the other player y , as shown in the right figure.

Fig. 4: Experimental process and task interface.

interface is shown in figure 4c. After participants completed all 15 sessions, the experiment ended.

C. Experimental Procedure

Prior to the experiment, participants were instructed not to engage in any interaction with each other. Initially, each participant signed a consent form and filled in a demographic survey. To familiarize themselves with the setup, two practice sessions were provided, wherein a practice drone was used to assist the participants. The participants were informed that the practice drone differed from the two drones used in the real experiment. After the experiment started, the assignment of drones was randomized for each pair of participants.

To motivate participants to provide accurate trust ratings, team performance instead of individual performance was used to determine the bonus, which was calculated as $\$10 \times$

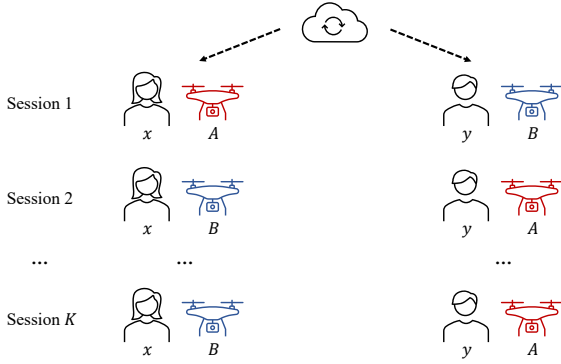


Fig. 5: Illustration of drone assignment. Participant x is randomly assigned to work with drone A in session 1, with drone B in session 2, and so on. The assignment is random.

$\max\{0, (\bar{a} - 0.7)/0.3\}$, where \bar{a} was the average detection accuracy of the two participants over all the tasks. Specifically, the participants would receive a bonus if their average detection accuracy exceeded 70%. Participants were explicitly informed that truthful and accurate communication of their trust values would assist the other participant in determining the appropriate level of trust in the drones, thereby increasing their detection accuracy and potential bonus.

V. RESULTS AND DISCUSSION

A. Analysis of Trust Convergence within Teams

We first conduct two types of team-level analysis to demonstrate that leveraging both direct and indirect interaction with a robot leads to faster trust convergence at the team level. We then compare the with- vs. between-team trust deviation and illustrate statistically the existence and benefits of leveraging both direct and indirect experience for trust updating. We denote the set of participants as $P = \{x_1, y_1, \dots, x_{15}, y_{15}\}$, where x_i and y_i are two members in the i th team.

Within-team trust average over time. We calculate the within-team trust average for team i on drone R at session k as

$$t_k^{i,R} := \frac{1}{2}(t_k^{x_i,R} + t_k^{y_i,R}),$$

where $R \in \{A, B\}$ indicates the drone type. The within-team trust average represents a team of players' overall trust in a robot.

Figure 6 shows how the within-team average trust changed as the number of interactions increased. The initial and final trusts in drone A ($\frac{1}{15} \sum_{i=1}^{15} t_0^{i,A}$ and $\frac{1}{15} \sum_{i=1}^{15} t_{15}^{i,A}$) were 0.57 ± 0.16 (mean \pm SD) and 0.83 ± 0.09 , respectively. The initial and final trusts in drone B ($\frac{1}{15} \sum_{i=1}^{15} t_0^{i,B}$ and $\frac{1}{15} \sum_{i=1}^{15} t_{15}^{i,B}$) were 0.61 ± 0.15 and 0.46 ± 0.19 , respectively. A two-way repeated measures analysis of variance (ANOVA) showed a significant main effect of drone type (drone A vs. B , $F(1, 14) = 58.81$, $p < .001$), and a non-significant effect of time (initial vs. final, $F(1, 14) = 3.66$, $p = .08$). There was also a significant interaction effect ($F(1, 14) = 73.02$, $p < .001$). Prior to the experiment, the within-team average trust in drone A and that




Fig. 6: Within-team trust *average* in drones A and B over time. Solid lines indicate mean values and the shaded regions indicate the sample standard errors.

in drone B were similar. As the amount of interaction increased, the within-team average trust in drones A and B tended to reflect the different detection accuracy of drone A and drone B , which were set to 90% and 60%, respectively. The within-team average trust in drone A gradually increased and that in drone B decreased. At the end of the experiment, the within-team average trust in drone A was significantly larger than that in drone B ($p < 0.001$).

Within-team trust deviation over time. We define the within-team trust deviation of team i on drone R at session k as the difference in trust ratings between the two human players in a team, regardless of whether the trust update is due to direct or indirect interaction, calculated as

$$\text{dev}_{k,W/N}^{i,R} := |t_k^{x_i,R} - t_k^{y_i,R}|,$$

where $R \in \{A, B\}$ is the drone type and the subscript “W/N” stands for “within.” In contrast to the within-team trust average, the within-team trust deviation focuses on the differences between the two players in a team.

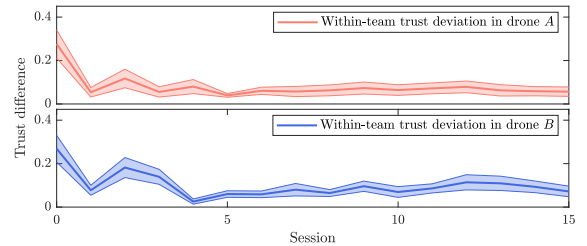


Fig. 7: Within-team trust *deviation* in drones A and B over time. Solid lines indicate the mean values and the shaded regions indicate the sample standard errors.

Figure 7 plots the within-team trust deviation in drone A and drone B . For both drones, the within-team trust deviation decreased rapidly in the first few sessions and became relatively stable afterward. For drone A , the initial and final within-team trust deviations ($\frac{1}{15} \sum_{i=1}^{15} \text{dev}_{0,W/N}^{i,A}$ and $\frac{1}{15} \sum_{i=1}^{15} \text{dev}_{15,W/N}^{i,A}$) were 0.27 ± 0.25 and 0.06 ± 0.08 . For drone B , the initial and final trust deviation values ($\frac{1}{15} \sum_{i=1}^{15} \text{dev}_{0,W/N}^{i,B}$ and $\frac{1}{15} \sum_{i=1}^{15} \text{dev}_{15,W/N}^{i,B}$) were 0.27 ± 0.24 and 0.07 ± 0.09 . A two-way repeated measures ANOVA revealed a significant main

effect of time, that the within-team trust deviation at the end of the experiment was significantly smaller than that prior to the experiment ($F(1, 14) = 11.51, p = .004$). Neither the drone type ($F(1, 14) = .06, p = .82$) nor the interaction effect ($F(1, 14) = .313, p = .59$) was significant.

Within- vs. between-team trust deviation. To statistically show the existence of trust propagation among team members, we compare the within-team and between-teams trust deviations as human agents gain more interaction experience. If trust propagation between the two players in a team had not occurred (i.e., participants updated their trust in the drones based solely on direct interaction), the within-team and between-team trust deviation would be statistically equal throughout the entire experiment. The between-team trust deviation of the i th team on drone R after the k th session is defined as

$$\text{dev}_{k,\text{BTW}}^{i,R} := \frac{1}{N-2} \sum_{p \in P \setminus \{x_i, y_i\}} \frac{1}{2} \left(|t_k^{x_i,R} - t_k^{p,R}| + |t_k^{y_i,R} - t_k^{p,R}| \right),$$

where $R \in \{A, B\}$ and N was the total number of participants. Figure 8 illustrates the calculation of within- and between-team trust deviations.




Fig. 8: Within-team trust deviation of team i is the trust difference between x_i and y_i , indicated by the dashed line in the figure. Between-team trust deviation of team i is the average trust difference between the trust ratings of x_i and y_i and all the other participants in other teams, indicated by the solid lines.

Figure 9 shows the within- vs. between-team trust deviations at the beginning and end of the experiment. In the beginning, the within- and between-team trust deviations in drone A were 0.28 ± 0.25 and 0.27 ± 0.22 , respectively, and in drone B were 0.27 ± 0.24 and 0.25 ± 0.21 , respectively (figure 9(a)). A two-way repeated measures ANOVA showed no significant difference between the within- and between-team trust deviation ($F(1, 14) = .07, p = .90$). No difference was found between drone A and drone B ($F(1, 14) = 2.82, p = .12$). The interaction effect was not significant either ($F(1, 14) = 0.75, p = .40$).

At the end of the experiment, the within-team and between-team trust deviations in drone A were 0.06 ± 0.08 and 0.11 ± 0.04 , and in drone B were 0.07 ± 0.09 and 0.22 ± 0.08 (figure

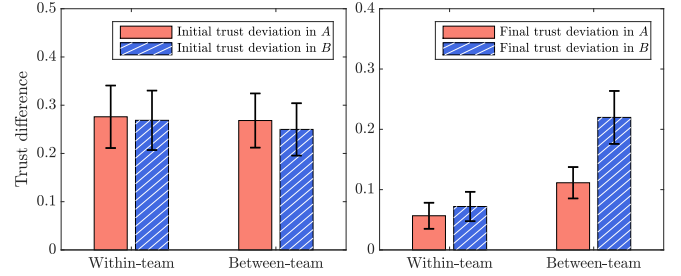


Fig. 9: Within- vs. between-team trust deviations (a) at the beginning and (b) end of the experiment

9(b)). A two-way repeated measures ANOVA revealed that the within-team trust deviation is significantly smaller than the between-team deviation ($F(1, 14) = 71.16, p < .001$), and trust deviation in drone A is significantly smaller than drone B ($F(1, 14) = 9.81, p = .007$). In addition, there was also a significant interaction effect ($F(1, 14) = 5.86, p = .03$).

The above results demonstrate the existence, and more importantly, the benefits of trust propagation. As shown in figures 6 and 7, the within-team trust average quickly stabilized and the within-team trust deviation rapidly decreased because of trust propagation within a team. Statistically speaking, at the beginning of the experiment, the within-team and between-team trust deviation in both drones were not significantly different (see figure 9(a)). At the end of the experiment, the within-team trust deviation was significantly smaller than the between-team trust deviation (see figure 9(b)). Had there not been trust propagation between the two players in a team (i.e., participants update their trust in the drones based *only* on the direct interaction), the within-team and between-team trust deviations would remain statistically equal. Therefore, the significant difference at the end of the experiment was attributed to the trust propagation within a team. Being able to fuse one's direct and indirect experience, instead of relying solely on the direct experience, contributes to the quick convergence of trust assessments on a robot, leading to a significantly smaller within-team trust deviation compared to the between-team trust deviation.

B. Model Fitting

For clarity, we relabel the participants as $P = \{p_1, p_2, \dots, p_{30}\}$. We utilize the gradient descent method in Section III-D to compute the optimal parameters $\Theta_*^{p_i,A}$ and $\Theta_*^{p_i,B}$ for each participant p_i . The fitting results are shown in figure 10. We set the performance measurements of drone A at session k as $p_k^A = A_k/10$ and $\bar{p}_k^A = 1 - p_k^A$, where A_k is the number of correct choices drone A made in the k th session; and we define p_k^B and \bar{p}_k^B similarly. To measure the performance of the model, we define the fitting error at each session for each participant as

$$e_k^{p_i,R} = \left| \mu_k^{p_i,R} - t_k^{p_i,R} \right|, \quad R \in \{A, B\},$$

where $t_k^{p_i,R}$ is the participant's reported trust while $\mu_k^{p_i,R}$ is the expected trust computed according to Eq. (2) with $\alpha_k^{p_i,R}$

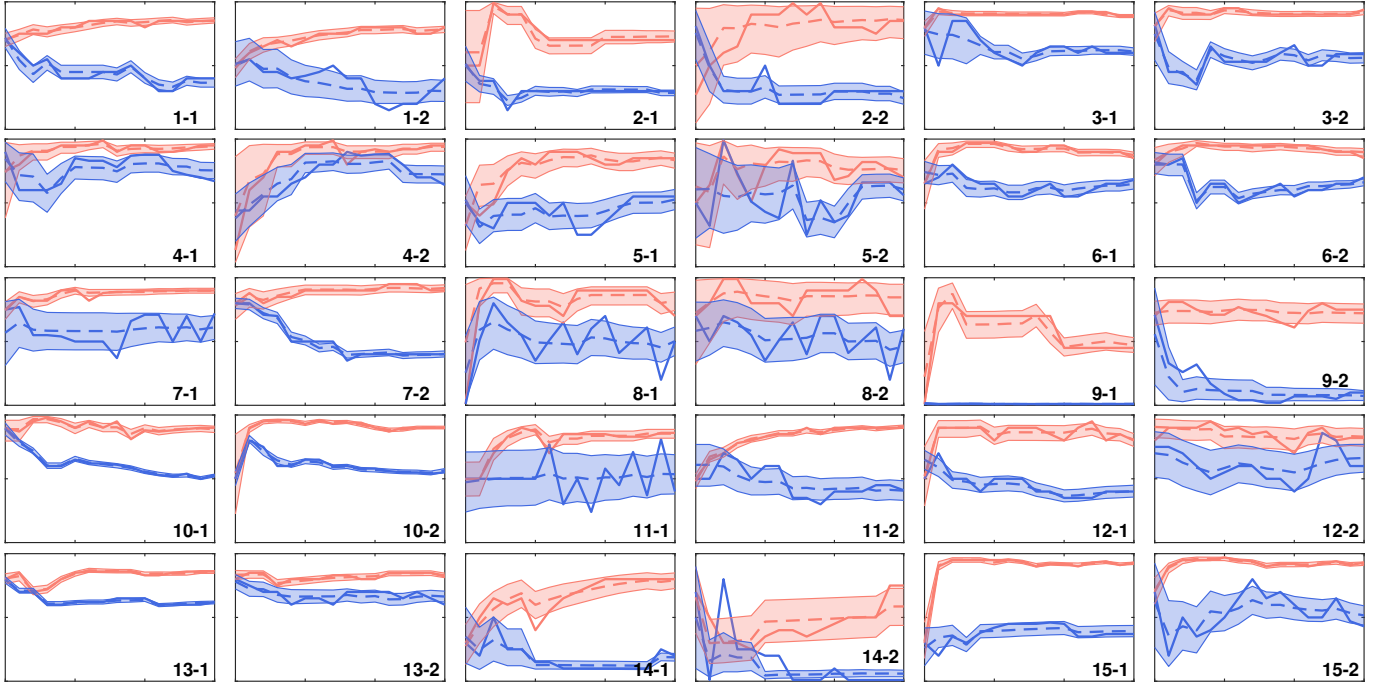


Fig. 10: Fitting results. Red curves are for drone A while blue curves are for drone B . The solid lines are the participants’ trust feedback, while the dashed lines are the expected trust value given by the model. The shaded areas indicate the 90% probability interval of the Beta distribution at each session. The index i - j stands for the j th participant in the i th group. The horizontal axes represent the session number, ranging from 0 (prior interaction index) to 15. The vertical axes indicate trust levels, ranging from 0 to 1, where 0 represents “(do) not trust at all” and 1 indicates “trust completely”.

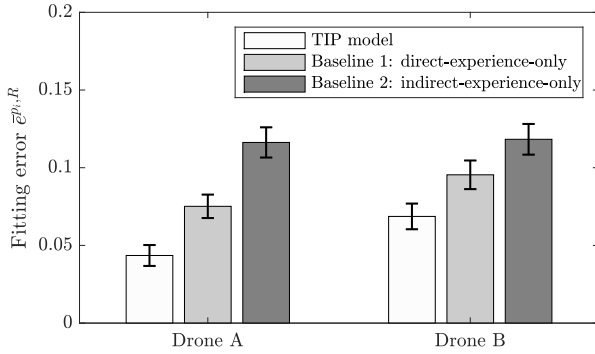


Fig. 11: Fitting error comparison between the TIP model and the two baseline models.

and $\beta_k^{p_i,R}$ generated by Eq. (8) based on $\theta_*^{p_i,R}$; and, we define the root-mean-square error (RMSE) between the ground truth and the expected trust value as

$$\text{RMSE}^R = \left[\frac{1}{N} \sum_{i=1}^N \frac{1}{K+1} \sum_{k=0}^K \left(e_k^{p_i,R} \right)^2 \right]^{1/2},$$

for $R \in \{A, B\}$. The results are $\text{RMSE}^A = 0.057$ and $\text{RMSE}^B = 0.082$.

For comparison, we consider two baseline models: one accounting for solely direct experience and another solely

indirect experience. The direct-experience-only model corresponds to the TIP model with zero unit gains in indirect experience, i.e., $\hat{s}^{x,A} = \hat{f}^{x,A} = 0$; while the indirect-experience-only model corresponds to $s^{a,b} = f^{a,b} = 0$. We recompute the parameters for the baseline models, and the RMSE errors are $\text{RMSE}_{\text{direct}}^A = 0.085$, $\text{RMSE}_{\text{direct}}^B = 0.107$, $\text{RMSE}_{\text{indirect}}^A = 0.128$, and $\text{RMSE}_{\text{indirect}}^B = 0.130$. In addition, we compare each participant’s fitting error $\bar{e}^{p_i,R} := 1/(K+1) \sum_{k=0}^K e_k^{p_i,R}$ of the TIP model (A : 0.044 ± 0.037 ; B : 0.069 ± 0.045), direct-experience-only model (A : 0.075 ± 0.041 ; B : 0.095 ± 0.051), and indirect-experience-only model (A : 0.116 ± 0.053 ; B : 0.118 ± 0.054) using a paired-sample t-test. Results reveal that the fitting error of the TIP model is significantly smaller than the direct-experience-only model, with $t(29) = -6.18, p < .001$ for drone A , and $t(29) = -7.31, p < .001$ for drone B , and significantly smaller than the indirect-experience-only model, with $t(29) = -9.28, p < .001$ for drone A , and $t(29) = -10.06, p < .001$ for drone B . Furthermore, the fitting error of the direct-only model is significantly smaller than the indirect-experience-only model, with $t(29) = -4.73, p < .001$ for drone A , and $t(29) = -3.73, p < .001$ for drone B . A bar plot is shown in figure 11. This comparison indicates that a human agent mainly relies on direct experience to update his or her trust, while indirect experience also plays a vital role in trust dynamics.

C. Trust Estimation

To measure the estimation accuracy of the proposed model, we remove some trust ratings in the data and compute the RMSE of the estimated trust values. Specifically, for each participant p_i , we set $U_{\hat{K}} = \{K - \hat{K} + 1, \dots, K\}$ to remove the last \hat{K} trust ratings, where $U_{\hat{K}}$ is the index set of sessions without trust ratings as defined in Section III-E, and estimate the missing trust values by $\mu_u^{p_i,A}$ and $\mu_u^{p_i,B}$ for $u \in U_{\hat{K}}$. The root-mean-square errors are defined as

$$\text{RMSE}_{\hat{K}}^R = \left[\frac{1}{N} \sum_{i=1}^N \frac{1}{\hat{K}} \sum_{u=K-\hat{K}+1}^K (e_u^{p_i,R})^2 \right]^{\frac{1}{2}},$$

for $R \in \{A, B\}$.

Figure 12 shows the RMSE's under different \hat{K} . When $\hat{K} \leq 7$, the TIP model can successfully estimate the trust values in the late sessions with a small RMSE (< 0.1) by learning from previous data. In particular, $\text{RMSE}_{\hat{K}=7}^A = 0.052$ and $\text{RMSE}_{\hat{K}=7}^B = 0.077$, which implies that, with the first 9 sessions' trust ratings available, the RMSE's of the estimation for the last 7 sessions are under 0.08 for both drones. The result also illustrates that $\text{RMSE}_{\hat{K}}^A$ is smaller than $\text{RMSE}_{\hat{K}}^B$ in general. This could be explained by the performance difference between the two drones. Indeed, because the number of correct choices each drone could make follows binomial distributions ($\text{Bin}(10, 0.9)$ for A and $\text{Bin}(10, 0.6)$ for B), the variance of their performance are 0.09 and 0.24 respectively. The greater variance of drone B may cause a human subject to acquire more information to stabilize his or her trust and thus leads to higher uncertainty in trust feedback values, which makes it difficult for the model to learn trust dynamics in a short time.

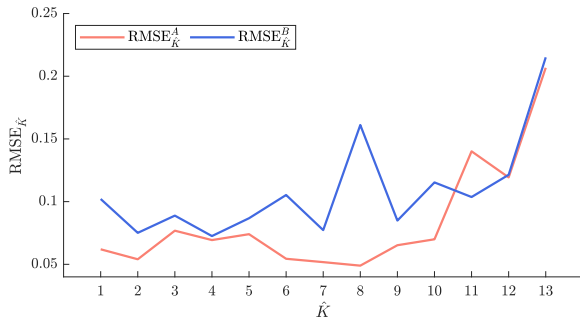


Fig. 12: RMSE's under different \hat{K} . \hat{K} is the number of interactions without trust feedback. Equivalently, the TIP model uses the trust ratings from the first $16 - \hat{K}$ sessions to learn the model parameters.

VI. CONCLUSION

In the study, we proposed the TIP model that accounts for both the direct and indirect experiences a human agent may have with a robot in multi-human multi-robot teams. To the best of our knowledge, it is the first mathematical framework for computational trust modeling in multi-human multi-robot

teams. In addition, we prove theoretically that trust converges after repeated direct and indirect interactions under our TIP framework. Using a human-subject experiment, we showed that being able to fuse one's direct and indirect experiences, instead of relying solely on the direct experience, contributes to the quick convergence of trust in a robot. In addition, we showed that the TIP models significantly outperformed the baseline direct-experience-only model in capturing the trust dynamics in multi-human multi-robot teams. The TIP model can be applied to various human-robot teaming contexts including team of teams [9] and multi-echelon networks [7]. In particular, the TIP model can update a human agent's trust in a robot whenever a direct or indirect experience is available and thus can be applied for trust estimation in a network consisting of multiple humans and robots.

Our results should be viewed in light of several limitations. First, we assume that the two human players within a team were cooperative and willing to share their trust in a robot truthfully. In a non-cooperative context where a human player is motivated to cheat, a quick convergence of trust assessment is less likely to occur. Further research is needed to examine the non-cooperative scenario. Second, we used a one-dimensional trust scale in the experiment. Even though this scale has been used in prior literature [10], [19], [20], it may not capture the different underlying dimensions of trust. Third, we take an ability/performance-centric view of trust and assume a human agent's trust in a robot is primarily driven by the ability or performance of the robot. Based on research in organizational management, trust can be influenced by three elements, namely ability, integrity, and benevolence [33]. Future research should investigate ways to integrate the benevolence and integrity elements into trust modeling, in particular, for HRI contexts that involve a strong emotional component, for example, educational or home-care robots. Moreover, conducting further ablation studies is essential to comprehensively understand the impact of various factors on the dynamics of trust. For instance, varying the number of sessions versus drone performances would provide insights into the rate at which trust converges.

REFERENCES

- [1] J. D. Lee and K. A. See, "Trust in automation: designing for appropriate reliance." *Human Factors*, vol. 46, no. 1, pp. 50–80, 2004.
- [2] T. B. Sheridan, "Human–Robot Interaction: Status and Challenges," *Human Factors*, vol. 58, no. 4, pp. 525–532, 2016.
- [3] X. J. Yang, C. Schemanske, and C. Searle, "Toward Quantifying Trust Dynamics: How People Adjust Their Trust After Moment-to-Moment Interaction With Automation," *Human Factors*, p. 00187208211034716, 2021.
- [4] H. Soh, S. Pan, C. Min, and D. Hsu, "The transfer of human trust in robot capabilities across tasks," in *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018.
- [5] R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human Factors*, vol. 39, no. 2, pp. 230–253, 1997.
- [6] P. Hancock, T. T. Kessler, A. D. Kaplan, J. C. Brill, and J. L. Szalma, "Evolving trust in robots: specification through sequential and comparative meta-analyses," *Human Factors*, vol. 63, no. 7, pp. 1196–1229, 2021.
- [7] National Academies of Sciences, Engineering, and Medicine, *Human-AI Teaming: State-of-the-Art and Research Needs*. Washington, DC: The National Academies Press, 2022.

- [8] W. F. Meehan and K. S. Jonker, "Team Of Teams: An Emerging Organizational Model," 2018. [Online]. Available: <https://www.forbes.com/sites/meehanjonker/2018/05/30/team-of-teams-an-emerging-organizational-model/?sh=321e2de36e79>
- [9] S. McChrystal, T. Collins, D. Silverman, and C. Fussell, *Team of Teams: New Rules of Engagement for a Complex World*. New York, New York: Penguin Publishing Group, 2015.
- [10] 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.
- [11] J. B. Lyons, T. Vo, K. T. Wynne, S. Mahoney, C. S. Nam, and D. Gallimore, "Trusting autonomous security robots: The role of reliability and stated social intent," *Human Factors*, vol. 63, no. 4, pp. 603–618, 2021.
- [12] M. Gombolay, X. J. Yang, B. Hayes, N. Seo, Z. Liu, S. Wadhwanina, T. Yu, N. Shah, T. Golen, and J. Shah, "Robotic assistance in the coordination of patient care," *The International Journal of Robotics Research*, vol. 37, no. 10, pp. 1300–1316, 2018, _eprint: <https://doi.org/10.1177/0278364918778344>.
- [13] N. Wang, D. V. Pynadath, and S. G. Hill, "Trust Calibration within a Human-Robot Team: Comparing Automatically Generated Explanations," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2016, pp. 109–116.
- [14] R. Luo, N. Du, and X. J. Yang, "Evaluating Effects of Enhanced Autonomy Transparency on Trust, Dependence, and Human-Autonomy Team Performance over Time," *International Journal of Human-Computer Interaction*, vol. 38, no. 18–20, pp. 1962–1971, 2022.
- [15] P. Robinette, W. Li, R. Allen, A. M. Howard, and A. R. Wagner, "Overtrust of Robots in Emergency Evacuation Scenarios," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. ACM, 2014, pp. 101–108.
- [16] E. J. de Visser, M. M. Peeters, M. F. Jung, S. Kohn, T. H. Shaw, R. Pak, and M. A. Neerinx, "Towards a theory of longitudinal trust calibration in human-robot teams," *International Journal of Social Robotics*, vol. 12, no. 2, pp. 459–478, 2020.
- [17] Y. Guo and X. J. Yang, "Modeling and Predicting Trust Dynamics in Human-Robot Teaming: A Bayesian Inference Approach," *International Journal of Social Robotics*, 2020.
- [18] J. D. Lee and N. Moray, "Trust, control strategies and allocation of function in human-machine systems," *Ergonomics*, vol. 35, no. 10, pp. 1243–1270, 1992.
- [19] D. Manzey, J. Reichenbach, and L. Onnasch, "Human performance consequences of automated decision aids: The impact of degree of automation and system experience," *Journal of Cognitive Engineering and Decision Making*, vol. 6, no. 1, pp. 57–87, 2012.
- [20] 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*, ser. HRI '17. New York, NY, USA: ACM, 2017, pp. 408–416.
- [21] X. J. Yang, Y. Guo, and C. Schemanske, "From Trust to Trust Dynamics: Combining Empirical and Computational Approaches to Model and Predict Trust Dynamics in Human-Autonomy Interaction," in *Human-Automation Interaction: Transportation*, V. G. Duffy, S. J. Landry, J. D. Lee, and N. A. Stanton, Eds., 2023, pp. 253–265.
- [22] 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*, 2018, pp. 307–315.
- [23] A. Xu and G. Dudek, "Optimo: Online probabilistic trust inference model for asymmetric human-robot collaborations," in *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2015, pp. 221–228.
- [24] N. Wang, D. V. Pynadath, and S. G. Hill, "The impact of pomdp-generated explanations on trust and performance in human-robot teams," in *Proceedings of the 2016 international conference on autonomous agents & multiagent systems*, 2016, pp. 997–1005.
- [25] H. Soh, Y. Xie, M. Chen, and D. Hsu, "Multi-task trust transfer for human-robot interaction," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 233–249, 2020.
- [26] B. C. Kok and H. Soh, "Trust in robots: Challenges and opportunities," *Current Robotics Reports*, vol. 1, no. 4, pp. 297–309, 2020.
- [27] C. Dellarocas, "The digitization of word of mouth: Promise and challenges of online feedback mechanisms," *Management Science*, vol. 49, no. 10, pp. 1407–1424, oct 2003.
- [28] F. Hendriks, K. Bubendorfer, and R. Chard, "Reputation systems: A survey and taxonomy," *Journal of Parallel and Distributed Computing*, vol. 75, pp. 184–197, 2015.
- [29] A. Jøsang, "Artificial reasoning with subjective logic," in *Proceedings of the second Australian workshop on commonsense reasoning*, vol. 48. Citeseer, 1997, p. 34.
- [30] A. Josang and R. Ismail, "The beta reputation system," in *Proceedings of the 15th bled electronic commerce conference*, vol. 5, 2002, pp. 2502–2511.
- [31] T. D. Huynh, N. R. Jennings, and N. Shadbolt, "Fire: An integrated trust and reputation model for open multi-agent systems," *Auton Agent Multi-Agent Syst*, vol. 13, p. 119–154, 2004.
- [32] Y. Guo, C. Shi, and X. J. Yang, "Reverse psychology in trust-aware human-robot interaction," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4851–4858, 2021.
- [33] R. C. Mayer, J. H. Davis, and F. D. Schoorman, "An integrative model of organizational trust," *Academy of management review*, vol. 20, no. 3, pp. 709–734, 1995.

APPENDIX

A. Proof of Theorem 1

Case 1: $n = 0$.

First, we show that $t_k^{x,A}$ converges i.p.. When $n = 0$, x gains experience towards A through only direct interaction, and thus, by Eq. (3),

$$\begin{aligned} \lim_{k \rightarrow \infty} \mu_k^{x,A} &= \lim_{k \rightarrow \infty} \frac{\alpha_k^{x,A}}{\alpha_k^{x,A} + \beta_k^{x,A}} \\ &= \lim_{k \rightarrow \infty} \frac{\alpha_0^{x,A} + k s^{x,A} r}{\alpha_0^{x,A} + \beta_0^{x,A} + k f^{x,A} \bar{r} + k s^{x,A} r} \\ &= \frac{s^{x,A} r}{f^{x,A} \bar{r} + s^{x,A} r}. \end{aligned} \quad (15)$$

For any $\epsilon > 0$, by the Markov inequality,

$$\begin{aligned} &\lim_{k \rightarrow \infty} \Pr\left(|t_k^{x,A} - \mu_k^{x,A}| < \epsilon\right) \\ &\leq \lim_{k \rightarrow \infty} \frac{1}{\epsilon^2} \mathbb{E}[(t_k^{x,A} - \mu_k^{x,A})^2] = 0, \end{aligned} \quad (16)$$

where the last equality is true because $\lim_{k \rightarrow \infty} \text{var}(t_k^{x,A}) = 0$ as $\lim_{k \rightarrow \infty} (\alpha_k^{x,A} + \beta_k^{x,A}) = \infty$. Let $t^x = \frac{s^{x,A} r}{f^{x,A} \bar{r} + s^{x,A} r}$. Eqs. (15) and (16) yield

$$\lim_{k \rightarrow \infty} \Pr\left(|t_k^{x,A} - t^x| < \epsilon\right) = 1.$$

Second, we show $|t_{k-1}^{y,A} - t_k^{x,A}|$ converge i.p. to zero. Suppose this is not true. Then, there exist $\epsilon, \delta > 0$ such that there are infinite many k 's such that $\Pr\left(|t_{k-1}^{y,A} - t_k^{x,A}| > \epsilon\right) > \delta$. The indirect updating rule Eq. (4) implies $\alpha_k^{y,A} + \beta_k^{y,A} \xrightarrow{\text{i.p.}} \infty$. Consequently, similar to Eq. (16), we obtain $|t_k^{y,A} - \mu_k^{y,A}| \xrightarrow{\text{i.p.}} 0$. We consider the following equation:

$$\begin{aligned} &\left| \mu_{k+m-1}^{y,A} - t_{k+m}^{x,A} \right| \\ &= \left| \mu_{k-1}^{y,A} - t_k^{x,A} \right| \prod_{j=1}^m \frac{\left| \mu_{k+j-1}^{y,A} - t_{k+j}^{x,A} \right|}{\left| \mu_{k+j-2}^{y,A} - t_{k+j-1}^{x,A} \right|}. \end{aligned} \quad (17)$$

As we have shown $t_k^{x,A} \xrightarrow{\text{i.p.}} t^x$ when $k \rightarrow \infty$, by the continuous mapping theorem, when $k \rightarrow \infty$,

$$\prod_{j=1}^m \frac{\left| \mu_{k+j-1}^{y,A} - t_{k+j}^{x,A} \right|}{\left| \mu_{k+j-2}^{y,A} - t_{k+j-1}^{x,A} \right|} \xrightarrow{\text{i.p.}} \prod_{j=1}^m \frac{\left| \mu_{k+j-1}^{y,A} - t^x \right|}{\left| \mu_{k+j-2}^{y,A} - t^x \right|}. \quad (18)$$

An examination of Eqs. (2) and (4) shows

$$\prod_{j=1}^m \frac{\left| \mu_{k+j-1}^{y,A} - t^x \right|}{\left| \mu_{k+j-2}^{y,A} - t^x \right|} \xrightarrow{\text{i.p.}} 0 \quad (19)$$

when $m \rightarrow \infty$. Eqs. (17), (18), and (19) together yield $|\mu_{k+m-1}^{y,A} - t_{k+m}^{x,A}| \xrightarrow{\text{i.p.}} 0$ when both k and m tend to infinity.

Thus, $\mu_k^{y,A} \xrightarrow{\text{i.p.}} t^x$. Because we also showed $|t_k^{y,A} - \mu_k^{y,A}| \xrightarrow{\text{i.p.}}$

0 , $|t_k^{y,A} - t^x| \xrightarrow{\text{i.p.}} 0$. This contradicts our assumption. Therefore, $|t_{k-1}^{y,A} - t_k^{x,A}|$ converge i.p. to zero. Particularly, since $t_k^{x,A} \xrightarrow{\text{i.p.}} t^x$, $t_k^{y,A} \xrightarrow{\text{i.p.}} t^x$.

Therefore, both $t_k^{x,A}$ and $t_k^{y,A}$ converge to t^x i.p..

Case 2: $n > 0$.

First, when $k \rightarrow \infty$, $\alpha_k^{x,A}$, $\beta_k^{x,A}$, $\alpha_k^{y,A}$, and $\beta_k^{y,A}$ all go to infinity as both x and y will have an infinite number of direct interactions with A . As a result, we have

$$\begin{aligned} &\left| t_k^{x,A} - \mu_k^{x,A} \right| \xrightarrow{\text{i.p.}} 0 \quad \text{and} \\ &\left| t_k^{y,A} - \mu_k^{y,A} \right| \xrightarrow{\text{i.p.}} 0. \end{aligned} \quad (20)$$

Second, let $\Delta t_k := t_k^{x,A} - t_k^{y,A}$. With the similar technique used in Eqs. (17), (18), and (19), it can be shown that

$$\Delta t_k \xrightarrow{\text{i.p.}} \Delta t, \quad (21)$$

where Δt is some constant.

Finally, we show $t_k^{x,A}$ and $t_k^{y,A}$ converge to constants i.p.. Since $|t_k^{x,A} - \mu_{k-1}^{x,A}| \xrightarrow{\text{i.p.}} 0$, by Eqs. (20) and (21), we obtain

$$\left| t_k^{y,A} - \mu_{k-1}^{x,A} \right| \xrightarrow{\text{i.p.}} \Delta t.$$

By Eqs. (3) and (4), both indirect and direct experience gains of y converge to some constants i.p. Therefore, the ratio $\frac{\alpha_k^{y,A}}{\alpha_k^{y,A} + \beta_k^{y,A}}$ converge to some constants i.p. respectively, i.e., there exist some t^y such that $t_k^{y,A} \xrightarrow{\text{i.p.}} t^y$. Similarly, there exists some t^x such that $t_k^{x,A} \xrightarrow{\text{i.p.}} t^x$. \square

B. Proof of Theorem 2

When $n = 0$, Eq. (5) yields $t^x = t^y = \frac{s^{x,A} r}{f^{x,A} \bar{r} + s^{x,A} r}$, which agrees with the proof of theorem 1.

Now we consider the case $n > 0$. Let t^x and t^y be the equilibrium in the statement. If $S^x/F^x - S^y/F^y \geq 0$, it can be shown that

$$\lim_{k \rightarrow \infty} \Pr\left(t_k^{x,A} - t_k^{y,A} \geq 0\right) = 1.$$

As a result, we have

$$\left[t_{k-1}^{y,A} - t_k^{x,A} \right]^+ \xrightarrow{\text{i.p.}} 0 \quad \text{and} \quad \left[t_k^{x,A} - t_{k-1}^{y,A} \right]^+ \xrightarrow{\text{i.p.}} t^x - t^y. \quad (22)$$

When $k \rightarrow \infty$, by Eqs. (3), (4), and (22), we have

$$\begin{aligned} &\alpha_{kl+l}^{x,A} - \alpha_{kl}^{x,A} \xrightarrow{\text{i.p.}} S_A^x \\ &\beta_{kl+l}^{x,A} - \beta_{kl}^{x,A} \xrightarrow{\text{i.p.}} \hat{F}_A^x (t^x - t^y) + F_A^x, \end{aligned}$$

where $l = m + n$. This implies the trust gains are constant every other l interactions. By Eq. (1),

$$t^x = \lim_{k \rightarrow \infty} \frac{\alpha_k^{x,A}}{\alpha_k^{x,A} + \beta_k^{x,A}} = \frac{S_A^x}{S_A^x + \hat{F}_A^x (t^x - t^y) + F_A^x},$$

which proves the first equation in Eq. (5). The second equation in Eq. (5) can be proved similarly.

The proof for the case when $S_A^x/F_A^x - S_A^y/F_A^y < 0$ is similar. \square

C. Computing t_A^x and t_A^y

We consider the case of Eq. (5). When $n = 0$, the solution is given by Eq. (15). Assume $n \neq 0$. Solving Eq. (5) directly results in two cubic equations of t^x and t^y respectively. An exact solution can be derived from these equations on $[0, 1]^2$ but the process can be tedious. A more practical method is to exploit the Newton's method to approximate t^x and t^y . For example, letting $z = 1 - y$, Eqs. in (5) give

$$\begin{aligned} \hat{F}^x x^2 + \hat{F}^x z x + (F^x + S^x - \hat{F}^x) x - S^x &= 0, \quad \text{and} \\ \hat{S}^y z^2 + \hat{S}^y z x + (S^y + F^y - \hat{S}^y) z - F^y &= 0. \end{aligned}$$

We can define $f_1(x, z)$ and $f_2(x, z)$ to be the left-hand sides of above equations and define $\mathbf{f} = (f_1, f_2)$. We have

$$\begin{aligned} \mathbf{f}'(x, z) \\ = \begin{pmatrix} 2\hat{F}^x x + \hat{F}^x z + F^x + S^x - \hat{F}^x & \hat{F}^x x \\ \hat{S}^y z & 2\hat{S}^y z + \hat{S}^y x + S^y + F^y - \hat{S}^y \end{pmatrix}. \end{aligned}$$

We solve

$$(x_{k+1}, z_{k+1})^T = (x_k, z_k)^T - (\mathbf{f}')^{-1} \mathbf{f}(x_k, z_k)$$

iteratively for x and z and obtain $t^x = x$ and $t^y = 1 - z$ as the equilibrium.