Assistance-Seeking in Human-Supervised Autonomy: Role of Trust and Secondary Task Engagement

Dong Hae Mangalindan and Vaibhav Srivastava

Abstract— **Using a dual-task paradigm, we explore how robot actions, performance, and the introduction of a secondary task influence human trust and engagement. In our study, a human supervisor simultaneously engages in a target-tracking task while supervising a mobile manipulator performing an object collection task. The robot can either autonomously collect the object or ask for human assistance. The human supervisor also has the choice to rely on or interrupt the robot. Using data from initial experiments, we model the dynamics of human trust and engagement using a linear dynamical system (LDS). Furthermore, we develop a human action model to define the probability of human reliance on the robot. Our model suggests that participants are more likely to interrupt the robot when their trust and engagement are low during high-complexity collection tasks. Using Model Predictive Control (MPC), we design an optimal assistance-seeking policy. Evaluation experiments demonstrate the superior performance of the MPC policy over the baseline policy for most participants.**

I. Introduction

With ongoing technological advancements, autonomous systems are increasingly deployed across industrial, commercial, healthcare, and agricultural sectors. These systems can handle hazardous or repetitive tasks, optimizing human resource allocation. However, despite their efficiency, these systems often require human supervision for complex tasks and in uncertain environments [1]–[3].

Trust in human-robot interaction is defined as "the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability" [4]. A lack of trust in autonomous systems can lead to their disuse, while excessive trust can result in misuse [5], [6]. Therefore, for effective human-robot collaboration, human trust in autonomous systems needs to be carefully calibrated.

Supervisors managing autonomous agents must balance their own tasks with overseeing these agents. For example, in an orchard, a human collects fruits while ensuring the robotic co-worker functions correctly, intervening only when necessary to avoid compromising their own tasks. The robot should also operate with minimal interference to the human's tasks. In this work, using a dual-task paradigm, we study how the robot's assistance-seeking affects human trust and secondary task engagement. We also investigate how robot actions and performance impact human secondary-task engagement and performance. We model the dynamics of human trust and secondary-task engagement and utilize it to design an optimal assistance-seeking policy.

This work was supported in part by the NSF award ECCS 2024649 and the ONR award N00014-22-1-2813.

D. H. Mangalindan and V. Srivastava are with the Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI 48824 e-mail: {mangalin, vaibhav}@egr.msu.edu

Various factors influence human trust, including robot reliability, workload, and information transparency [1], [7]–[11]. These works highlight that negative experiences lead to a decrease in human trust. Additionally, a higher level of information transparency can have a positive effect on human trust but can increase their workload [1], [10]. Factors affecting human trust in automation are categorized into three main groups: robot-related, human-related, and environment-related factors, with robot-related factors such as reliability influencing the trust the most, followed by environment-related factors such as task complexity [12].

Dynamic models of trust consider these factors as inputs or states. Probabilistic trust models, such as Partially Observable Markov Decision Process (POMDP)-based models, treat the trust as a latent state and estimate its distribution within a Bayesian framework [13]. These models have been used for planning to minimize human interruption [11], [14], [15] and to optimize information provided to humans, maintaining collaboration efficiency [1].

Linear dynamical systems use trust level, cumulative trust level, and expectation bias as system states. The inputs to these models include factors that influence human trust, such as robot performance and task-specific contextual information [9], [10], [13]. Process and measurement noises have also been incorporated into linear models of human trust [16]. We focus on a similar noisy linear trust model; however, our measurement model is nonlinear (binary), which in contrast to [16], necessitates nonlinear filtering techniques, specifically the particle filter. Additionally, we leverage our learned model to determine and experimentally evaluate an optimal assistance-seeking policy.

In our dual-task paradigm, participants simultaneously perform supervisory and target-tracking tasks. We explore an assistance-seeking policy for the robot based on human trust, tracking task engagement, and task-specific contexts. In the initial data collection experiment, the robot uses a randomized assistance-seeking policy with varying task complexities. We use the collected data to model trust dynamics and task engagement using an LDS with robot and human actions, task outcomes, and task complexities as inputs. Additionally, we develop a model to predict human actions based on trust and engagement. Using MPC [17] with the estimated models, we compute robot actions to maximize team performance. We then evaluate the MPC-based policy in a second set of experiments, comparing it with a baseline policy. Our contributions include:

- Studying the effect of assistance-seeking on human trust and secondary-task performance;
- Designing an optimal-assistance-seeking policy based

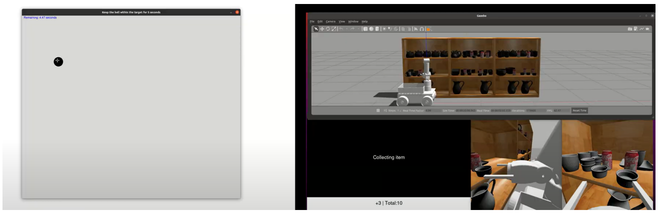

Fig. 1: Experiment interface showing the target-tracking and supervisory task rendered on two adjacent screens.The setup uses ROS-Gazebo and resources available in [18], [19].

on human behavior models;

- Evaluating the designed policy in experiments with human participants; and
- Using a nonlinear filter for real-time estimation of trust and secondary-task engagement.

The paper is organized as follows. Sections II and III describe the dual-task setup and the human behavior model, respectively. Section IV details the computation and evaluation of the optimal policy. Section V discusses our results and their broader implications. We conclude in Section VI.

II. Dual Task Setup: Robot Supervision and target-tracking task

In our experiment, the human participant engages in a dual-task paradigm, simultaneously supervising an autonomous mobile manipulator retrieving objects from shelves and performing a target-tracking task. The experiment interface is shown in Fig. 1.

A. Human Supervised Robotic Object Collection Task

In our supervisory task, the human supervisor collaborates with a robot to gather items from grocery shelves. The supervisory goal is to monitor the mobile manipulator and ensure its safety and efficiency. The human supervisor monitors a live feed displaying the robot's actions and views from two cameras placed on the end-effector (see Fig. 1; right panel).

In each trial, the mobile manipulator has the choice of attempting autonomous collection (a^{R+}) or requesting human assistance (a^{R-}), wherein the human must teleoperate the manipulator using a joystick and retrieve the item. When the robot initiates autonomous collection of an object, the human supervisor can either rely (a^{H+}) on the robot to collect the object autonomously, or intervene (a^{H-}) and collect the object via teleoperation.

The complexity (C^1) of each object collection task is determined by the object's location relative to the mobile manipulator. Specifically, the complexity is high ($C^1 = C^{1,\mathrm{H}}$) if the manipulator's direct path to the object is obstructed by an obstacle; otherwise, the complexity is low ($C^1 = C^{1,\mathrm{L}}$).

The outcome of each trial of the object collection task is either a success (E^+), if the object is collected successfully and dropped into the bin; or a failure (E^-), otherwise. The human-robot team reward for the object collection task is defined by

$$\mathcal{R}^{\text{coll}}(a^R, a^H, E) = \begin{cases} +3, & \text{if } (a^R, a^H, E) = (a^{R+}, a^{H+}, E^+), \\ +1, & \text{if } a^R = a^{H+}, \\ 0, & \text{if } (a^R, a^H) = (a^{R+}, a^{H-}), \\ -4, & \text{if } (a^R, a^H, E) = (a^{R+}, a^{H+}, E^-). \end{cases}$$

The objective of the human-robot team for the object collection tasks is to maximize the cumulative team reward across all trials.

B. Target-tracking Task

In addition to the supervisory task, the human participant also performs a target-tracking task (see Fig. 1; left panel). In each target-tracking trial, the human participant must maintain the computer cursor inside a moving ball using the mouse for a total duration of 5 seconds. In our experiment, we select two possible speeds (slow and normal) for the ball and they result in complexity $C^2 \in \{C^{2,\text{slow}}, C^{2,\text{norm}}\}$.

The human participants are instructed to track the ball (target) for a total cumulative time of 5 seconds. Their target-tracking performance p is determined by taking the goal time of 5 seconds and dividing it by the total time it took to finish the trial. The participants were instructed to achieve a performance of at least 0.75. Then, the reward for the target-tracking task is defined by

$$\mathcal{R}^{\text{track}}(C^2, p) = \begin{cases} +0.5, & \text{if } C^{2,\text{norm}} \text{ and } p \geq 0.75, \\ +0.25, & \text{if } C^{2,\text{slow}} \text{ and } p \geq 0.75, \\ 0, & \text{if } p < 0.75. \end{cases}$$

C. Data Collection for Modeling Human Behavior

To build a model of human supervisory behavior, we conducted a study[1] with 11 human participants (4 males and 7 females), performing the above supervisory and target-tracking tasks simultaneously. Participants were recruited via e-mail for the in-person experiment. Participants were first and second-year college students not in the engineering program. An experiment's total duration was one hour and participants were compensated $15 for their participation.

For the supervisory object collection tasks, each participant performed 30 low-complexity and 30 high-complexity trials, randomly distributed throughout the experiment. When the manipulator operates autonomously, it has a success probability of $p_H^{\text{suc}} = 0.75$ in $C^{1,\mathrm{H}}$ and $p_L^{\text{suc}} = 0.96$ in $C^{1,\mathrm{L}}$. In each trial, the robot asks for human assistance with a probability of 0.3 in $C^{1,\mathrm{H}}$ trials, and with a probability of 0.1 in $C^{1,\mathrm{L}}$ trials. After each trial, participants reported their trust in the autonomous mobile manipulator through an interface with an 11-point scale. For each object collection trial, the data collected are the complexity C^1, robot action a^R, human action a^H, outcome E, and the reported human trust y^T.

For the target-tracking task, participants performed 30 slow-speed trials and 30 normal-speed trials. To ensure that all participants do the same number of trials per complexity, an initial set of 60 such trials was selected and trials were randomly permuted for each participant and independently of the object-collection task. In each trial, the data collected are the complexity of the trial C^2, and the performance p.

D. Data Summary

For the supervisory task, after a successful autonomous collection, the mean trust ratings of participants are 7.38

[1]The human behavioral experiments were approved under Michigan State University Institutional Review Board Study ID 9452.

for low complexity and 7.36 for high complexity, with an overall mean trust rating of 7.37 across both complexities. After a failure, the mean trust ratings were 6.36 for low complexity and 6.68 for high complexity, with an overall mean of 6.6 across both complexities. When the robot asked for assistance, the mean trust ratings were 7.3 for low complexity and 6.98 for high complexity, averaging 7.03 for both complexities combined. However, when the human voluntarily interrupted, the mean trust rating dropped to 6.29 when they interrupted the robot in high-complexity, 5 in low-complexity, and 6.17 across both complexities.

For the tracking task, at slow speeds, the mean performance was 0.89 with 90% of the trials having $p \geq 0.75$. At normal speeds, the mean performance dropped to 0.82 with 78% of the trials having $p \geq 0.75$. As expected, $p < 0.75$ whenever participants interrupted the robot.

III. Human Behavior Model

In this section, we focus on a model of human supervisory behavior, i.e., a model that predicts when the human supervisor may intervene in the autonomous operation of the manipulator. We assume that human action is modulated by their trust in the mobile manipulator and their engagement in the target tracking task.

In the following, we denote trust at the beginning of trial t by $T_t \in \mathbb{R}$, the target-tracking engagement in trial $(t-1)$ by $G_t \in [0,1]$, the complexities of object collection and target-tracking task by C_t^1 and C_t^2, respectively, the human and robot actions during collaborative object collection in trial t by a_t^H and a_t^R, respectively, and the outcome of the object collection in trial t by E_{t+1}.

A. Human Trust Dynamics

The evolution of human trust in the robot is known to be influenced by several factors such as the current trust level, robot performance, robot actions, environmental complexity, and human actions. Accordingly, we assume that trust T_{t+1} at the beginning of trial $t+1$ is influenced by the previous trust T_t, robot action a_t^R, collection task complexity C_t^1, human action a_t^H, and the outcome of the collection task E_{t+1}.

Consider the following events corresponding to 6 different combinations of C^1, a^R, and E; and the additional case when human interrupts the robot corresponding to $(a_t^R, a_t^H) = (a^{R+}, a^{H-})$. For brevity, we define events by their outcomes.

θ_t^1	$(C_t^{1,L}, a_t^{R+}, a_t^{H+}, E_{t+1}^+)$	θ_t^5	$(C_t^{1,H}, a_t^{R+}, a_t^{H+}, E_{t+1}^-)$
θ_t^2	$(C_t^{1,L}, a_t^{R+}, a_t^{H+}, E_{t+1}^-)$	θ_t^6	$(C_t^{1,H}, a_t^{R-})$
θ_t^3	$(C_t^{1,L}, a_t^{R-})$	θ_t^7	(a_t^{R+}, a_t^{H-})
θ_t^4	$(C_t^{1,H}, a_t^{R+}, a_t^{H+}, E_{t+1}^+)$		

Let $u_t^T \in \{0,1\}^7$ be an indicator vector representing one of the above 7 possible events in collaborative object collection trial t. We model the trust dynamics using the following LDS

$$T_{t+1} = A^T T_t + B^T u_t^T + v_t^T, \tag{1}$$

where $A^T \in (0,1)$, $B^T \in [-1,1]^{1\times 7}$, and $v_t^T \sim \mathcal{N}(0, \sigma_T^2), t \in \mathbb{N}$, are i.i.d. realizations of zero-mean Gaussian noise. The entries of B^T determine the influence of the events θ_t^i on the evolution of trust.

We assume that the trust reported by the participants after each trial is possibly a scaled and noisy measurement of the trust T_t, i.e.,

$$y_t^T = C^T T_t + w_t^T, \tag{2}$$

where $C^T \in [0,1]$ and $w_t^T \sim \mathcal{N}(0, \sigma_y^2), t \in \mathbb{N}$, are i.i.d. realizations of zero-mean Gaussian noise.

The trust model parameters are $A^T, B^T, C^T, \sigma_T^2$ and σ_y^2. We adopt the Expectation-Maximization (EM) algorithm [20] with the reported trust values y_t^T to obtain the following estimates

$$A^T = 0.92, \quad \sigma_T^2 = 0.22, \quad C^T = 1.00, \quad \sigma_y^2 = 0.22,$$
$$B^T = \begin{bmatrix} 0.76 & -0.38 & 0.26 & 0.78 & -0.43 & 0.52 & -0.12 \end{bmatrix}.$$

The elements of B^T represent the effect of each scenario on human trust. The estimated values for contexts associated with a negative experience, e.g., a failed collection or human intervention are negative. Similarly, the estimated values for contexts associated with a positive experience, e.g., successful collections or when the robot requests assistance, are positive. Interestingly, as found in our previous study [14], asking for assistance can help increase human trust, which is also consistent with the estimated B^T. In other words, the estimates suggest that successful collection increases and maintains trust; failed collections decrease it; and asking for assistance can help repair and increase trust. Additionally, human interruptions decrease their trust.

B. Human Target-tracking Engagement Dynamics

Similar to modeling human trust dynamics, we aim to model how robot action a_t^R, tracking task complexity C_t^2, and previous experience in the supervisory task E_t influence human target-tracking engagement and performance. The experience E_t is the outcome of the autonomous collection trial; while, for engagement dynamics, we classify it as E^- if the human interrupts the robot, and as E^+, if the robot asks for assistance.

Consider the following events corresponding to 8 different combinations of C^2, a^R, and E. Note that for brevity, we have defined the events with their outcomes.

ϕ_t^1	$(C_t^{2,\text{slow}}, a_t^{R-}, E_t^+)$	ϕ_t^5	$(C_t^{2,\text{slow}}, a_t^{R+}, E_t^+)$
ϕ_t^2	$(C_t^{2,\text{norm}}, a_t^{R-}, E_t^+)$	ϕ_t^6	$(C_t^{2,\text{norm}}, a_t^{R+}, E_t^+)$
ϕ_t^3	$(C_t^{2,\text{slow}}, a_t^{R-}, E_t^-)$	ϕ_t^7	$(C_t^{2,\text{slow}}, a_t^{R+}, E_t^-)$
ϕ_t^4	$(C_t^{2,\text{norm}}, a_t^{R-}, , E_t^-)$	ϕ_t^8	$(C_t^{2,\text{norm}}, a_t^{R+}, E_t^-)$

Let $u_t^G \in \{0,1\}^8$ be an indicator vector representing one of the above 8 possible events in the target-tracking task during trial t. We model the human target-tracking engagement dynamics as an LDS

$$G_{t+1} = A^G G_t + B^G u_t^G + v_t^G \tag{3}$$
$$p_t = C^G G_t + w_t^G, \tag{4}$$

where $A^G \in [0,1]$, $B^G \in [0,10]^{1\times 8}$, $v_t^G \sim \mathcal{N}(0, \sigma_G^2), t \in \mathbb{N}$ are i.i.d. realizations of zero-mean Gaussian noise, $C^G \in [0,1]$, and $w_t^G \sim \mathcal{N}(0, \sigma_p^2), t \in \mathbb{N}$, are i.i.d. realizations of zero-mean Gaussian noise. The entries of B^G determine the influence of the events ϕ_t^i on the human engagement in the tracking task.

The engagement model parameters include $A^G, B^G, C^G, \sigma_G^2$ and σ_p^2. Considering p as the tracking task result of the human, we adopt the EM algorithm to estimate these parameters. The estimated parameters of the engagement dynamics model are

$$A^G = 0.19, \quad \sigma_G^2 = 1.44, \quad C^G = 9.96, \quad \sigma_p^2 = 3.79,$$
$$B^G = \begin{bmatrix} 7.47 & 6.72 & 7.24 & 6.38 & 7.30 & 6.51 & 7.06 & 6.59 \end{bmatrix}.$$

The elements of B^G represent the effect of each scenario on human tracking-task engagement and performance.

The estimated B^G indicates that whenever the robot asks for assistance, for the majority of cases, the associated estimated values are greater than the values associated with the cases when the robot collects autonomously. This is expected as whenever the robot asks for assistance, the human supervisor can focus more on their tracking task and finish it before teleoperating the robot, knowing that the robot has stopped and will not move unless they take control of it. In contrast, when the robot attempts to collect autonomously, the human supervisor may switch attention from the tracking task to the supervisory task to decide whether to rely on the robot. The estimated parameters are higher when the target speed is slow compared to when it is normal. Whenever the previous experience with the robot is E^-, the parameters are lower compared to the case of E^+. This is consistent with the fact that whenever humans perceive a failure, they tend to pay more attention to the supervisory task in the future.

C. Human Action Model

We now focus on modeling the probability of the human relying on the robot when the robot attempts to collect autonomously. We assume that this probability depends on the object collection task complexity C_t^1, human trust T_t, and human target-tracking engagement G_t, which is representative of the attention they allocate to the tracking task. We model the human action probability as a sigmoid function

$$\mathbb{P}(a^{H+}|T, G, C^1) = \frac{1}{1 + e^{-(a_T^{C^1} T + a_G^{C^1} G + b^{C^1})}}. \quad (5)$$

The model parameters are learned with Maximum Likelihood Estimation using the Monte Carlo simulation method, resulting in the following parameter values for different object collection task complexities

	$a_T^{C^1}$	$a_G^{C^1}$	b^{C^1}
$C^{1,L}$	0.09	0.08	3.6
$C^{1,H}$	0.20	0.40	-2.7

The probabilities of reliance for each object collection task complexity are shown in Fig. 2. For low complexity $C^{1,L}$, the probability of reliance is high and close to 1, regardless of T and G. For high complexity $C^{1,H}$, there is a lower probability of reliance in general as compared to low complexity $C^{1,L}$. It can be seen that the probability of reliance increases with trust as well as with tracking task engagement. The latter may be caused by the human focusing more on the tracking task as compared to the supervisory task, thereby increasing their probability of reliance on the mobile manipulator.

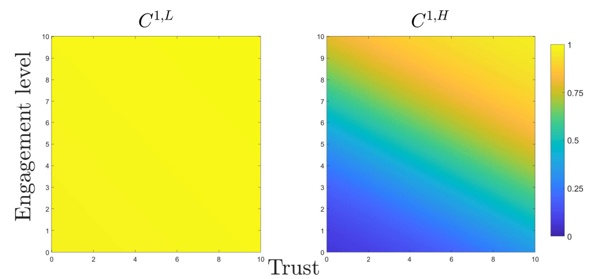

Fig. 2: Human action model: Probability of reliance $\mathbb{P}(a^{H+}|T, G, C^1, a^{R+})$ conditioned on object collection task complexity C^1,T, and F

IV. Design of Human Assistance Seeking Policy

In this section, we propose an optimal assistance-seeking policy, leveraging the estimated trust dynamics, target-tracking dynamics, and human action model.

A. MPC-based Assistance-seeking Policy

Using the models described in Section III, we design an optimal assistance-seeking policy using MPC. We adopt a certainty-equivalent MPC approach [17] that considers the expected reward and deterministic system dynamics. Specifically, the following optimization problem is solved at the beginning of each trial t with trust T_t and target-tracking engagement G_t.

$$\begin{aligned} \underset{q_t, \dots q_{t+N-1}}{\text{maximize}} \quad & \sum_{\tau=t}^{t+N-1} \bar{\mathcal{R}}^{\text{coll}}(\bar{T}_\tau, \bar{G}_\tau, q_\tau) + \bar{\mathcal{R}}^{\text{track}}(\bar{G}\tau, q_\tau) \\ \text{s.t.} \quad & \bar{T}_{\tau+1} = A^T \bar{T}_\tau + \hat{B}_\tau^T u_\tau, \; \bar{T}_t = T_t \\ & \bar{G}_{\tau+1} = A^G \bar{G}_\tau + \hat{B}_\tau^G u_\tau \;, \bar{G}_t = G_t \\ & \bar{p}_t = C^G \bar{G}_t \\ & q_\tau \in [0, 1], \end{aligned} \quad (6)$$

for $\tau = t, \dots, (t+N-1)$, where $u_\tau = \begin{bmatrix} q_\tau & 1 - q_\tau \end{bmatrix}^\top$, and q is the probability of taking action a^{R+}.

The entries in $\hat{B}_t^T$ and $\hat{B}_t^G$ correspond to the expected influence of the autonomous collection attempt and assistance-seeking on trust and engagement, respectively. They are obtained by averaging entries of B^T and B^G, respectively, associated with these events weighted by their probabilities. Similarly, $\bar{\mathcal{R}}^{\text{coll}}(\bar{T}_\tau, \bar{G}_\tau, q_\tau)$ and $\bar{\mathcal{R}}^{\text{track}}(\bar{G}\tau, q_\tau)$ correspond to the expected reward. The interested reader may refer to [21] for detailed calculations.

In the MPC framework, the optimization problem (6) is solved at the beginning of each trial t, the optimal sequence of robot actions $\{\bar{q}_t, \dots \bar{q}_{t+N-1}\}$ is computed and only the first action is executed, i.e., a_t^R is chosen as the autonomous collection a^{R+} with probability $\bar{q}_t$, and seeking assistance, otherwise. The look-ahead horizon N is a tuning parameter that is selected to balance the performance and computational time and to account for the accuracy of the look-ahead

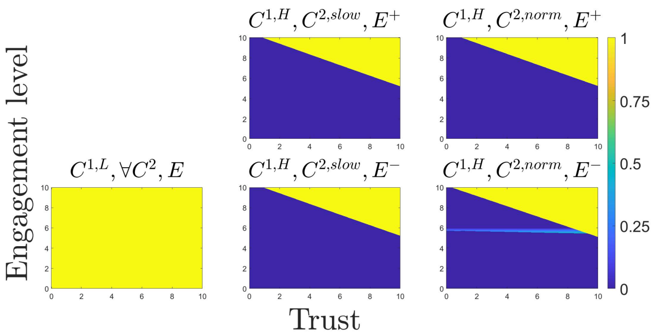

Fig. 3: MPC policy: The optimal action for $C^{1,\mathrm{L}}$ is to collect, whereas for $C^{1,\mathrm{H}}$, action switches based on the trust and engagement states.

model. We solved the optimization problem (6), using MATLAB's `fmincon` function. The assistance-seeking policy obtained using the above MPC formulation with a look-ahead horizon $N = 5$ is shown in Fig. 3.

In low-complexity object collection tasks, the optimal action for the robot is always to collect autonomously, regardless of other variables. This outcome is anticipated because of the low interruption rate and high success rate observed in low-complexity object collection tasks.

In high-complexity object collection tasks, the computed optimal action is to seek assistance if tracking task engagement falls below $G_t = 5$, regardless of other variables. Since low target-tracking performance may be caused by the human paying more attention to the supervisory task, the optimal policy seeks assistance at low target-tracking engagement; thereby allowing the human to finish the tracking task and then assist the manipulator. In contrast, when tracking task engagement is high, the optimal policy depends on human trust: it asks for assistance at low trust levels and collects autonomously for high trust levels. The threshold of trust for seeking assistance decreases with target-tracking performance. Since higher target-tracking performance is an outcome of higher engagement with the tracking task resulting in fewer interruptions by the supervisor, the optimal policy attempts autonomous collection of the object.

Interestingly, for high-complexity object collection and normal speed tracking tasks, the threshold of trust for asking assistance in the optimal policy is lower when the previous experience is negative (E^-) compared to when it is positive (E^+). Since the probability of interruption increases following a failure, the policy aims to mitigate this interruption by requesting assistance, which results in human trust repair and enables human to improve their tracking performance.

B. Evaluation Experiment

We conducted a second set of human experiments to evaluate the designed assistance-seeking policy and compare it to a baseline policy that only considers object-collection complexity and tracking speed. Using data from the initial experiment, we estimated the probability of reliance $\mathbb{P}(a^{H+}|C^1, C^2)$ on the robot and the probability of success $\mathbb{P}(p \geq 0.75|C^1, C^2)$ in the tracking task. We computed the expected reward for each robot action based on these probabilities. The optimal baseline policy, referred to as the "greedy policy", always attempts autonomous collection.

We recruited 5 participants (3 females and 2 males) who completed two experiment blocks where the robot followed the MPC and greedy policies, in a randomized order. Each block included 15 trials each of low- and high-complexity object collection tasks. To match a realistic scenario where participants might not report their trust after each trial, we estimated trust and task engagement based on their actions (rely a^{H+} or intervene a^{H-}) and measured tracking performance (p_t) instead of using reported trust values. To this end, we adopted a particle filter with trust dynamics (1), engagement dynamics (3), action model (5) and target-tracking performance output (4) to estimate T_{t+1} and G_t.

The cumulative reward for all participants under both the greedy and the MPC policy is illustrated in Fig. 4. The score of a trial is calculated as the sum of rewards in tracking and collection tasks. The median scores are 65.75 and 57 for the MPC and the greedy policy, respectively. The MPC policy outperforms the greedy policy for most participants. The cumulative number of interruptions in our MPC policy, 16, is also less than the greedy policy, 23.

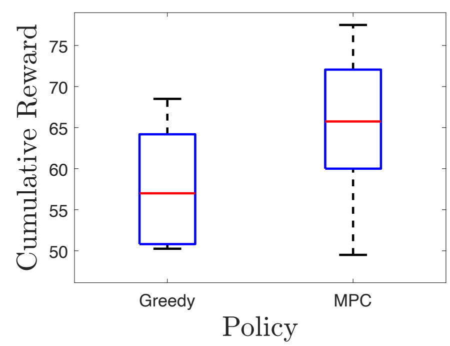

Fig. 4: Cumulative reward (sum of collection and tracking rewards) statistics for both policies

V. Discussion

When humans encounter a multitasking scenario, they must optimize their resources to achieve acceptable performance in every task. In the context of a human-robot team, human trust in the robot and their engagement in their own responsibilities play a key role in their multi-tasking performance. To understand and leverage human trust and engagement dynamics toward designing better policies for the robot, we developed a model and estimated it using human participant experiment data.

In our prior study [14], we investigated the supervisory object collection task without any secondary task. In contrast to the dual-task paradigm, we observed a higher interruption rate in our prior study (0.173 vs. 0.095). This is possibly due to the division of engagement across two tasks. Additionally, the optimal assistance-seeking policy in the single-task paradigm, as described by [14], seeks assistance only in cases of high task complexity and low human trust. In this paper, the robot's actions are guided by both the level of human trust and engagement in a secondary task. When engagement with the secondary task is low, resembling a single-task paradigm, the optimal policy consistently seeks assistance. The optimal policy considers the secondary task performance, which tends to decline at low engagement levels, and seeking assistance allowing humans to satisfactorily complete the secondary task before assisting.

Our model parameters indicate that failures decrease trust, aligning with the literature that negative experiences negatively affect trust [8], [10]. Similarly, when humans interrupt and intervene, their trust in the robot tends to decrease [11]. Our human action model suggests that humans are unlikely to rely on robots in scenarios associated with high complexity and low trust. This is consistent with the observations in the literature that at a higher risk, a higher trust is required for humans to rely on the robot [1], [11].

In a dual-task setup, when trust in the robot is low, they tend to interrupt the robot, reducing their attention to the other task and thereby diminishing their performance. When a human supervises an autonomous robot, it has been established that as human trust towards the robot increases, their rate of monitoring and interruption decreases [15]. When performing a dual task, if performance on one task diminishes, it is due to the other task consuming most of the operator's resources [22]. Therefore, we can consider their engagement in the tracking task as a secondary measure of their trust. When their trust in the robot is low, they tend to supervise it more closely, diverting attention from their tracking task, thus decreasing their engagement and diminishing their performance.

VI. Conclusions

We focused on a scenario in which a human supervises object collection by a robot while performing a secondary target-tracking task. We studied the impact of robot performance on the primary task and human performance on the secondary task on the evolution of trust and secondary task engagement. We modeled the trust and secondary task performance dynamics using linear dynamical systems with Gaussian noise and human action selection probability as a static function of their trust and secondary task engagement. A data-collection experiment with human participants was conducted and the collected data was used to estimate these models. We formulated an optimal assistance-seeking problem for the robot that seeks to optimize team performance while accounting for human trust and secondary task engagement and solved it using MPC. We showed that the optimal assistance-seeking policy is to never seek assistance in low-complexity object-collection trials and to seek assistance in high-complexity object-collection trials only when the trust is below a secondary-task-engagement-dependent threshold. Specifically, this threshold decreases with the engagement. We compared this proposed policy with a greedy baseline policy and conducted experiments with human participants to evaluate the MPC policy. The results showed that the MPC policy outperformed the greedy policy.

References

[1] K. Akash, G. McMahon, T. Reid, and N. Jain, "Human trust-based feedback control: Dynamically varying automation transparency to optimize human-machine interactions," *IEEE Control Systems Magazine*, vol. 40, no. 6, pp. 98–116, 2020.

[2] J. R. Peters, V. Srivastava, G. S. Taylor, A. Surana, M. P. Eckstein, and F. Bullo, "Human supervisory control of robotic teams: Integrating cognitive modeling with engineering design," *IEEE Control Systems Magazine*, vol. 35, no. 6, pp. 57–80, 2015.

[3] P. Gupta and V. Srivastava, "Structural properties of optimal fidelity selection policies for human-in-the-loop queues," *Automatica*, vol. 159, p. 111388, 2024.

[4] J. Lee and K. See, "Trust in technology: Designing for appropriate reliance," *Human Factors*, vol. 46, no. 1, pp. 50–80, 2004.

[5] K. A. Hoff and M. Bashir, "Trust in automation: Integrating empirical evidence on factors that influence trust," *Human Factors*, vol. 57, no. 3, pp. 407–434, 2015.

[6] R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human Factors*, vol. 39, no. 2, pp. 230–253, 1997.

[7] J. Lee and N. Moray, "Trust, control strategies and allocation of function in human-machine systems," *Ergonomics*, vol. 35, no. 10, pp. 1243–1270, 1992.

[8] M. Desai, P. Kaniarasu, M. Medvedev, A. Steinfeld, and H. Yanco, "Impact of robot failures and feedback on real-time trust," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2013, pp. 251–258.

[9] K. Akash, W.-L. Hu, T. Reid, and N. Jain, "Dynamic modeling of trust in human-machine interactions," in *2017 American Control Conference (ACC)*, 2017, pp. 1542–1548.

[10] 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 *ACM/IEEE International Conference on Human-Robot Interaction*, 2017, pp. 408–416.

[11] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa, "Trust-aware decision making for human-robot collaboration: Model learning and planning," *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 9, no. 2, pp. 1–23, 2020.

[12] P. A. Hancock, D. R. Billings, K. E. Schaefer, J. Y. Chen, E. J. De Visser, and R. Parasuraman, "A meta-analysis of factors affecting trust in human-robot interaction," *Human Factors*, vol. 53, no. 5, pp. 517–527, 2011.

[13] Y. Wang, F. Li, H. Zheng, L. Jiang, M. F. Mahani, and Z. Liao, "Human trust in robots: A survey on trust models and their controls/robotics applications," *IEEE Open Journal of Control Systems*, vol. 3, pp. 58–86, 2024.

[14] D. H. Mangalindan, E. Rovira, and V. Srivastava, "On trust-aware assistance-seeking in human-supervised autonomy," in *2023 American Control Conference (ACC)*. IEEE, 2023, pp. 3901–3906.

[15] Z. Zahedi, M. Verma, S. Sreedharan, and S. Kambhampati, "Trust-aware planning: Modeling trust evolution in iterated human-robot interaction," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2023, pp. 281–289.

[16] H. Azevedo-Sa, S. K. Jayaraman, C. T. Esterwood, X. J. Yang, L. P. Robert Jr, and D. M. Tilbury, "Real-time estimation of drivers' trust in automated driving systems," *International Journal of Social Robotics*, vol. 13, no. 8, pp. 1911–1927, 2021.

[17] J. Mattingley, Y. Wang, and S. Boyd, "Receding horizon control," *IEEE Control Systems Magazine*, vol. 31, no. 3, pp. 52–65, 2011.

[18] L. Downs, A. Francis, N. Koenig, B. Kinman, R. Hickman, K. Reymann, T. B. McHugh, and V. Vanhoucke, "Google scanned objects: A high-quality dataset of 3D scanned household items," in *International Conference on Robotics and Automation*, 2022, pp. 2553–2560.

[19] G. Ramkumar and J. Lentin, *ROS Robotics Projects : Build and Control Robots Powered by the Robot Operating System, Machine Learning, and Virtual Reality.* Packt Publishing, 2019, vol. Second edition.

[20] K. P. Murphy, *Probabilistic Machine Learning: Advanced Topics*. MIT Press, 2023. [Online]. Available: http://probml.github.io/book2

[21] D. H. Mangalindan and V. Srivastava, "Assistance-seeking in human-supervised autonomy: Role of trust and secondary task engagement (extended version)," *arXiv preprint arXiv:2405.20118*, 2024.

[22] S. M. Casner and B. F. Gore, "Measuring and evaluating workload: A primer," *NASA Technical Memorandum*, vol. 216395, p. 2010, 2010.