

Clustering Trust Dynamics in a Human-Robot Sequential Decision-Making Task

Shreyas Bhat¹, Joseph B. Lyons², Cong Shi¹, and X. Jessie Yang¹

Abstract—In this paper, we present a framework for trust-aware sequential decision-making in a human-robot team. We model the problem as a finite-horizon Markov Decision Process with a reward-based performance metric, allowing the robotic agent to make trust-aware recommendations. Results of a human-subject experiment show that the proposed trust update model is able to accurately capture the human agent's moment-to-moment trust changes. Moreover, we cluster the participants' trust dynamics into three categories, namely, Bayesian decision makers, oscillators, and disbelievers, and identify personal characteristics that could be used to predict which type of trust dynamics a person will belong to. We find that the disbelievers are less extroverted, less agreeable, and have lower expectations toward the robotic agent, compared to the Bayesian decision makers and oscillators. The oscillators are significantly more frustrated than the Bayesian decision makers.

I. INTRODUCTION

Trust has been identified as a key factor for effective human-robot interaction. Consequently, substantial research efforts have been devoted to identifying factors that influence humans' trust in robots [1], developing computational models for trust estimation [2], [3], and developing trust-aware robots [4], [5].

Previous computational models of trust [2], [3] require binary performance measures of the autonomy. However, such measures cannot be directly applied in scenarios wherein the robotic agent needs to perform complex trade off decisions to maximize the cumulative reward and hence the autonomy performance is more difficult to quantify. In addition, although previous studies have revealed the existence of different types of trust dynamics [3], [6], no research has tried to associate personal characteristics with the type of trust dynamics.

To fill these research gaps, we develop a novel reward-based performance metric to drive the trust estimation algorithm in a Markov Decision Process (MDP). The proposed performance metric is a binary surrogate of the autonomous agent's reward function, which allows our autonomous agent to utilize the trust update model developed in [3] for trust-aware decision-making. Moreover, from the data collected through human-subject experiments, we analyze how human trust evolves with their earned reward over repeated

interactions with the autonomy. We find three distinct types of trust dynamics through k -means clustering analysis and examine associations between personal characteristics and type of trust dynamics.

II. RELATED WORK

Extensive work has been done in identifying factors that affect trust in automation. However, most of these studies provide a snapshot view of trust, and failed to characterize the dynamics of trust [7]. More recently, dynamic models of trust has been proposed, including the auto-regressive moving average vector (ARMAV) model [8], the Online Probabilistic Trust Inference Model (OPTIMO) model [2], and the experience-based model [3]. These models enable a robot to estimate a human's trust in real time and has led to the development of trust-aware robots, allowing robots to reason upon human trust in their planning [4], [5], [9], [10].

Most existing research investigating individual differences aims to find associations between (snapshot) trust and individual characteristics, including propensity to trust automation [11], personality trait of neuroticism [12], and dispositional expectations of technology [13], [14]. It is possible that these individual differences will be differentially represented within trust profiles related to trust dynamics. Recent research [3] reveal the existence of different types of trust dynamics, namely, rational decision maker, oscillator, and disbeliever. It is reasonable to investigate the association between a human agent's trust dynamics and factors that are shown to affect a human agent's (snapshot) trust.

III. PROBLEM FORMULATION

In this section, we propose a finite horizon Markov Decision Process (MDP) for modeling and incorporating trust in the decision-making system of a robotic agent. The robot provides recommendations to its human partner about the action that they should take, but the final decision lies with the human. The specific scenario we target is an 'Intelligence, Surveillance, and Reconnaissance' (ISR) mission in which a human soldier teams up with an *intelligent* drone to search through a town for threats. The drone guides the soldier on whether s/he should breach a site directly or deploy a Robotic Armored Rescue Vehicle (RARV). Using the RARV prevents any health loss to the soldier in the presence of a threat but takes additional time to the team; on the contrary, breaching a site directly is faster, but the soldier will be harmed if s/he encounters a threat. Here, two natural (but conflicting) goals that arise are to minimize any damage to the soldier and minimize the time to search through all the sites.

*This work was supported in part by the Air Force Office of Scientific Research (Grant #FA9550-20-1-0406).

¹Shreyas Bhat, Cong Shi, and X. Jessie Yang are with the Industrial and Operations Engineering Department, University of Michigan, Ann Arbor, MI 48109 (email: shreyasb@umich.edu, shicong@umich.edu, xijyang@umich.edu)

²Joseph B. Lyons is with the Air Force Research Laboratory (email: joseph.lyons.6@us.af.mil)

The key components of the MDP model of the ISR mission are given below.

1) **States:** We use the estimated trust of the human on the robot as the state. Specifically, we use the trust model in [3], wherein trust level at the i^{th} stage, t_i , follows a Beta distribution, i.e., $t_i \sim \text{Beta}(\alpha_i, \beta_i)$, where α_i and β_i are the positive experience and negative experience at time i . And, the state at stage i is specified by the tuple (α_i, β_i) .

2) **Actions:** The drone has two actions: Recommend to use or not use the RARV.

3) **Human behavior model:** We assume that human will accept the recommendation given by the robotic agent with probability t_i ; and the human chooses the opposite action of the one that was recommended with probability $1 - t_i$.

4) **Reward function:** We use the weighted sum of the health loss cost and the time cost to define the cost of choosing an action. Similar to [5], we add a trust-gaining term to incentivize the autonomous agent to gain trust. As a result, the reward R_i at the i^{th} stage is defined as

$$R_i(a) = -w_h h(a) - w_c c(a) + \gamma_i \cdot \mathbb{1}(A), \quad (1)$$

where w_h and w_c are the weights for the health loss cost and time loss cost respectively. We define A as the event when trust increases, i.e., $\mathbb{1}(A) = 1$ if the performance of the autonomous agent is a success and $\mathbb{1}(A) = 0$ otherwise. $\gamma_i (= w_t \sqrt{N-i})$ is a weight given to the trust gain reward that decreases with the stage number. The idea behind this is to support trust-gaining behavior near the current stage, and task performance optimizing behavior towards the later stages of planning.

5) **Transition function:** The state (α_i, β_i) is updated by

$$(\alpha_i, \beta_i) = \begin{cases} (\alpha_{i-1} + w^s, \beta_{i-1}) & \text{if } p_i = 1, \\ (\alpha_{i-1}, \beta_{i-1} + w^f) & \text{if } p_i = 0. \end{cases} \quad (2)$$

Here, p_i is the autonomous agent's performance at the i^{th} stage. Since the team's goal is to maximize earned reward, we define $p_i = 1$ if at the i^{th} site, the immediate reward for following the recommendation by the drone was greater than that for not following the recommendation and $p_i = 0$ otherwise.

We fit the parameters $(\alpha_0, \beta_0, w^s, w^f)$ for each participant to model their trust dynamics. We use gradient descent to estimate the parameters in real time after each feedback from the participant. We use value iteration to solve the MDP.

IV. EXPERIMENT

This section describes details of the human-subject experiment. The experiment complied with the American Psychological Association code of ethics and was approved by the Institutional Review Board at the University of Michigan.

A. Participants

A total of 46 adults participated in the study. One participant's data was discarded as the participant failed the attention check. The remaining 45 participants consisted of 21 females and 24 males (Age: Mean = 22.8 years, $SD = 3.6$

years). Each participant was reimbursed with a base pay of \$20 with a bonus of up to \$10 based on their performance, which was measured by the time taken by the participants to complete the task and the final health level of the soldier.

B. Testbed

We developed a 3D testbed in Unreal Engine. A screenshot of the testbed is shown in Fig. 1(a). Fig. 1(b) shows the recommendation dialog box where the participant was recommended to not use the RARV. After choosing an action, the four possibilities depending on the presence of threat and participant's selected actions are shown in Fig. 2. The participants are told that if they encounter a threat without the RARV, they will lose 5 points of health. They were told that deploying the RARV takes about 10 seconds. They were told to choose an action based on their interaction history and the recommendation. Their goal was to minimize the time taken for the mission, while maintaining the soldier's health.




Fig. 1. The testbed developed in the Unreal Engine 4 game engine

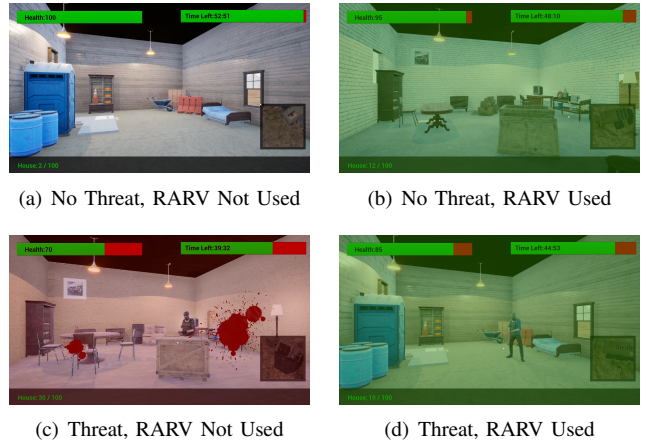


Fig. 2. The outcomes based on threat presence and action chosen.

C. Measures

1) **Personality:** The big 5 factors of personality (Extraversion, Conscientiousness, Agreeableness, Neuroticism, and Imagination) were measured using the 20-item mini-IPIP scale [15]. This 5-point Likert scale has widely been used in human-robot trust research [16].

2) **Propensity to Trust:** Propensity to trust autonomous systems was assessed using a 6-item scale developed in [17]. This was also a 5-point Likert scale

3) *Perfect Automation Schema*: Perfect Automation Schema (PAS) was measured using the 7-item scale developed by [13]. Of these, 4 items measure high expectations from the autonomy and the other 3 items measure All-or-none thinking. This was a 7-point Likert scale.

4) *Moment-to-Moment Trust*: During the task, the participants were asked to rate their moment-to-moment trust on the drone by adjusting a slider on a 100-point scale.

5) *Post-experiment trust*: After the experiment, we used two scales [16], [18] for assessing trust. The first one was a 9-item questionnaire with sliders while we used 6-items from the second scale all of which were 7-point Likert type questions.

6) *Workload*: Workload was measured using the NASA Task Load Index [19]. We only used 5 of the 6 items as there was no physical demand from the participants in our experiment. All the items had the participants rate their feelings on a slider with values ranging from Very low to Very high.

D. Experimental procedure

Prior to the experiment, participants provided informed consent and completed the pre-experiment surveys. They were oriented to the steps of the experiment and walked through each of the screens they would see during the experiment. The two-fold objective of minimizing time and maximizing health was emphasized. They were informed about the performance-based bonus pay. They were told that the drone was imperfect, but they were not informed of the exact reliability level. They were also told that the robotic agent's recommendations would help them achieve the two-fold objective. Participants then proceeded to the experimental trials, wherein they had to search through 100 houses sequentially. After searching each house, the participants were asked to report their level of trust on the autonomous agent's recommendations. The participants took on average 51.5 minutes to complete the task. At the end of the experiment, participants completed the post-experiment surveys.

V. RESULTS & DISCUSSION

A. Using Immediate Actual Reward as a performance metric

Since the participants are explicitly told to consider both the soldier's health and the time to complete the search as their objectives, we expect their trust to be correlated with the immediate reward that they receive upon choosing an action. We expect that a participant's trust would be likely to increase if following the recommendation by the drone gets them a higher reward than doing the opposite and vice versa. Fig. 3 shows a representative example. In the figure, a green triangle represents the site at which following the recommendation would result in a better immediate reward gain ($p_i = 1$) and the red triangles represent the opposite. It is quite clear that a red triangle is often followed by a decrease in trust and a green triangle is followed by an increase in trust. Thus, our reward-based performance metric is able to capture moment-to-moment trust changes of the




Fig. 3. Trust feedback with overlay of red and green triangles representing the rewards.

participant. Using this metric in our trust update model, we get a prediction root mean squared error of 0.1266 ($SD = 0.078$).

B. Clustering of Trust Dynamics




Fig. 4. Clustering of participants according to their trust dynamics.

We employ k-means clustering to group participants with similar trust dynamics. We use the root mean squared error (E_{RMS}) between the predicted trust and the trust feedback and the average logarithm of trust as the clustering features. For computing E_{RMS} , we consider the feedback at the first 20 sites as a training set to fit the personalized parameters ($\alpha_0, \beta_0, w^s, w^f$). Thereafter, we use the feedback after every 5 sites to update the parameters. Fig. 4 shows the results. We choose $k = 3$ as the optimum number of clusters because it is a turning point of both the within-cluster variance and the silhouette score. We call the cluster with small E_{RMS} and generally higher values of trust *Bayesian Decision Makers*. The second cluster, called *Disbelievers*, consists of participants whose E_{RMS} values are small but whose trust is generally low. The third significant group is the *Oscillators*, whose trust changes rapidly, making it harder to predict.

C. Personal Characteristics and Type of Trust Dynamics

We performed one-way ANOVA between the data of the clustered participants. Results showed significant difference between the three types of trust dynamics in Extraversion ($F(2,42) = 4.991, p = 0.011$), Agreeableness ($F(2,42) =$

3.276, $p = 0.048$), and the High Expectations facet of the Perfect Automation Schema ($F(2, 42) = 5.752, p = 0.006$). Further, propensity to trust automation ($F(2, 42) = 3.002, p = 0.06$) and intellect/imagination ($F(2, 42) = 2.687, p = 0.08$) had marginally significant differences.

Post-hoc analysis with Bonferroni adjustment shows that the disbelievers are significantly less extroverted than the oscillators ($p = 0.009$). Disbelievers were only marginally less extroverted than Bayesian Decision Makers ($p = 0.061$). Disbelievers were marginally less agreeable than the Bayesian Decision Makers ($p = 0.068$) and Oscillators ($p = 0.06$). Disbelievers had significantly lower expectations from automation compared to both Bayesian Decision Makers ($p = 0.005$) and Oscillators ($p = 0.023$).

One way ANOVA on the post-experiment measures show significant difference between the three clusters in their post-experiment trust reports (Trust questionnaire by Muir and Moray $F(2, 42) = 22.167, p < 0.001$, Trust questionnaire by Lyons and Guznov $F(2, 42) = 15.183, p < 0.001$) and their frustration levels ($F(2, 42) = 4.136, p = 0.023$).

Post-hoc analysis with Bonferroni adjustment shows that there are significant differences between each of the three groups' trust reports according to the trust questionnaire by Muir and Moray ($p < 0.001$ between Bayesian Decision Makers and Disbelievers, $p = 0.006$ between Bayesian Decision Makers and Oscillators, and $p = 0.009$ between Oscillators and Disbelievers with the highest trust for Bayesian Decision Makers and lowest for Disbelievers). Trust reported with the questionnaire by Lyons and Guznov only showed significant difference between Disbelievers and Bayesian Decision Makers ($p < 0.001$), and between Oscillators and Bayesian Decision Makers ($p = 0.002$). Oscillators were significantly more frustrated than Bayesian Decision Makers ($p = 0.025$).

VI. CONCLUSIONS

In this study, we formulated the human-robot sequential decision-making problem as an MDP and proposed an innovative reward-based performance metric for trust-aware decision-making. Through a human-subject experiment, we found three distinct types of trust dynamics, namely, Bayesian Decision Makers, Oscillators, and Disbelievers. Our model achieved different accuracies on different types on trust dynamics, which points to a requirement to use different models for people in different categories.

Analyses suggest that those individuals classified as Bayesian Decision makers evidenced high expectations of automation. When combined with other state measures such as trust and frustration, one might be able to identify a profile of the Bayesian Decision makers as the combination of individual differences and state measures could be used to parse the sample into the three categories of trust dynamics. Knowing that an individual might fall into one of the categories could influence whether or not a machine partner that is equipped with a dynamic trust model is a feasible solution for that individual.

In future studies, besides the reverse-psychology model, different human trust-behavior models should be considered. In addition, techniques such as inverse reinforcement learning can be employed to learn the preference of the human teammate to improve the team performance.

REFERENCES

- [1] P. A. 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, 2020.
- [2] 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.
- [3] Y. Guo and X. J. Yang, "Modeling and predicting trust dynamics in human-robot teaming: A bayesian inference approach," *International Journal of Social Robotics*, 12 2021.
- [4] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa, "Trust-aware decision making for human-robot collaboration: Model learning and planning," *J. Hum.-Robot Interact.*, vol. 9, no. 2, jan 2020. [Online]. Available: <https://doi.org/10.1145/3359616>
- [5] 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.
- [6] G. McMahon, K. Akash, T. Reid, and N. Jain, "On modeling human trust in automation: Identifying distinct dynamics through clustering of markovian models," *IFAC-PapersOnLine*, vol. 53, pp. 356–363, 2020.
- [7] 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.
- [8] 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.
- [9] A. Xu and G. Dudek, "Maintaining efficient collaboration with trust-seeking robots," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 3312–3319.
- [10] K. Akash, T. Reid, and N. Jain, "Improving human-machine collaboration through transparency-based feedback – part ii: Control design and synthesis," *IFAC-PapersOnLine*, vol. 51, no. 34, pp. 322–328, 2019, 2nd IFAC Conference on Cyber-Physical and Human Systems CPHS 2018.
- [11] S. Merritt and D. Ilgen, "Not all trust is created equal: Dispositional and history-based trust in human-automation interactions," *Human factors*, vol. 50, pp. 194–210, 05 2008.
- [12] J. L. Szalma and G. S. Taylor, "Individual differences in response to automation: the five factor model of personality," *Journal of experimental psychology. Applied*, vol. 17 2, pp. 71–96, 2011.
- [13] S. M. Merritt, J. L. Unnerstall, D. Lee, and K. Huber, "Measuring individual differences in the perfect automation schema," *Human Factors*, vol. 57, no. 5, pp. 740–753, 2015.
- [14] J. Lyons and S. Guznov, "Individual differences in human-machine trust: A multi-study look at the perfect automation schema," *Theoretical Issues in Ergonomics Science*, vol. 20, pp. 1–19, 11 2018.
- [15] M. B. Donnellan, F. L. Oswald, B. M. Baird, and R. E. Lucas, "The mini-ipp scales: tiny-yet-effective measures of the big five factors of personality," *Psychological assessment*, vol. 18, no. 2, p. 192, 2006.
- [16] J. B. Lyons, C. S. Nam, S. A. Jessup, T. Q. Vo, and K. T. Wynne, "The role of individual differences as predictors of trust in autonomous security robots," in *2020 IEEE International Conference on Human-Machine Systems (ICHMS)*, 2020, pp. 1–5.
- [17] S. M. Merritt, H. Heimbaugh, J. LaChapell, and D. Lee, "I trust it, but i don't know why: Effects of implicit attitudes toward automation on trust in an automated system," *Human Factors*, vol. 55, no. 3, pp. 520–534, 2013, pMID: 23829027. [Online]. Available: <https://doi.org/10.1177/0018720812465081>
- [18] B. Muir and N. Moray, "Trust in automation. part ii. experimental studies of trust and human intervention in a process control simulation," *Ergonomics*, vol. 39, pp. 429–60, 04 1996.
- [19] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Human Mental Workload*, ser. Advances in Psychology, P. A. Hancock and N. Meshkati, Eds. North-Holland, 1988, vol. 52, pp. 139–183.