Towards Efficient Collaborations with Trust-Seeking Adaptive Robots

Anqi Xu
McGill University
Montreal, Canada
anqixu@cim.mcgill.ca

Gregory Dudek
McGill University
Montreal, Canada
dudek@cim.mcgill.ca

ABSTRACT

We are interested in asymmetric human-robot teams, where a human supervisor occasionally takes over control to aid an autonomous robot in a given task. Our research aims to optimize team efficiency by improving the robot's task performance, decreasing the human's workload, and building trust in the team. We envision synergistic collaborations where the robot adapts its behaviors dynamically to optimize efficacy, reduce manual interventions, and actively seek for greater trust. We describe recent works that study two facets of this *trust-seeking adaptive* methodology: modeling human-robot trust dynamics, and developing interactive behavior adaptation techniques. We also highlight ongoing efforts to combine these works, which will enable future human-robot teams to be maximally trusting and efficient.

Figure 1: Our interactive, adaptive vehicle tracking a road. Inset: camera view overlaid with tracked boundary (blue line), robot's steering (blue arrow), and human's intervening command (green arrow).

1. INTRODUCTION

Trust – one's belief in another's competence and reliability – is the cornerstone of all long-lasting collaborations, both among human teammates, as well as between humans and robots. In the latter context, the degree of trust that a human operator has in an autonomous robot is strongly correlated to the team's performance, and also greatly influences the operator's behaviors [5, 3, 7]. Teams harboring high levels of trust often demonstrate efficiency and synergies, where the human and robot work together to complement each other's skills and weaknesses. In contrast, low trust can degenerate into disuse of automation, where the human stops delegating tasks to the robot, or disables it altogether. In extreme cases, such distrust has led to fatal accidents, for instance in train derailments where the automated alert systems were disabled due to prior false alerts [6].

The objective of our research is to improve the efficiency of human-robot teams. We quantify *team efficiency* as a combination of performance metrics, such as task error and automation failures, as well as human factors, including workload, satisfaction, and trust. We assume that the autonomous robot is always motivated and obedient; thus the human's trust in the robot arises due to *performance-based* causes only, as opposed to concerns of deception or defiance.

Our studies focus on navigation tasks where the human supervises a mobile robot as it visually tracks terrain boundaries. Concretely, we built aerial robots that can follow highways and shorelines, and wheeled robots that learn to drive along roads and trails (Fig. 1). The human operator in these teams occasionally intervenes and assumes control, either to teach the robot to track a new target, or to correct misbehaviors. Visual navigation tasks are appealing since humans innately excel at them, whereas the required complexity in autonomous solutions [3, 7] warrants the need for trust.

Towards our goal of efficient human-robot teams, we propose to develop *trust-seeking adaptive robots*. These robots will be able to sense when the human has low trust, and reactively adapt their behaviors to improve performance and seek greater trust. We present two recent research achievements: a dynamic *human-robot trust model*, and an *online behavior adaptation* technique. We also discuss ongoing efforts to combine these works and reach our objective of long-lasting and high-trusting human-robot teams.

2. HUMAN-ROBOT TRUST MODEL

Before the robot can elicit greater trust, it must first be able to reason about the human's trust state. We developed a near real-time model that can accurately predict the operator's moment-to-moment trust states, ranged between distrust and full trust [8]. This is achieved by inferring beliefs about trust during interactions through experience factors that are known to correlate strongly to trust [5, 3, 7].

Our model is captured by a Dynamic Bayesian Network, and updates the trust belief as a linear function of the robot's

HRI'15 Extended Abstracts, March 2–5, 2015, Portland, OR, USA.
ACM 978-1-4503-3318-4/15/03.
http://dx.doi.org/.

current and recent task performance (i.e. causal attribution for trustworthiness), as well as a bias term (i.e. propensity to award trust). These probabilistic trust estimates are then calibrated to match the human's latest intervention states, where the likelihood of a user taking manual control is modeled as a linear logistic expression dependent on a constant bias (i.e. predisposition to micromanage), low trust and loss of trust (i.e. trust-related causes), and extraneous causes to intervene (e.g. training to follow a new boundary target).

We conducted a controlled study with 20 roboticists to collect interaction datasets for boundary patrol tasks, within a simulated environment. In addition to logging task performance and intervention states, we occasionally queried the users' *absolute trust* feedback, through an interval scale. We also asked users during interactions to periodically report *trust changes*, by pressing gamepad buttons indicating trust lost, gained, or unchanged. Our Bayesian model incorporates these variable-rate trust assessments to further ground beliefs about the human's latent trust state.

We trained instances of this dynamic trust model on half of each user's dataset using Expectation-Maximization. Trust predictions made by these personalized models for the hold-out portion of each dataset consistently outperformed those of several existing works [4, 3, 7], both in terms of numerical error and Pearson's r correlation. This improvement arises from our model's use of a temporal probabilistic representation, which captures the *uncertainty* in trust inferences across time. Another uniqueness of our model is the ability to infer trust states rapidly and every few seconds, whereas prior models operated at coarser time scales of minutes or longer, due to their regression formulation. The high accuracy and near real-time contributions of this model will empower next-generation, responsive, and trust-aware robots.

3. ONLINE BEHAVIOR ADAPTATION

Our research also investigates methods to adapt and improve the robot's behaviors online, both in autonomous and manual control states. Focusing on the latter, we developed an interactive adaptation technique that learns from human interventions to improve team efficiency [9]. This is based on policy gradient reinforcement learning [2] and resembles Learning from Demonstration methods [1]. We generalized the framework to handle *changing task goals*, and realized an online, anytime, and concurrent solution.

The robot's policy $y_r = \mathbb{A}(x, \theta)$ computes control outputs y_r (e.g. steering commands) given sensor states x (e.g. camera frames), and configurable parameters θ. Human commands y_h can be used to compare intervening behaviors against simulated policy outputs for arbitrary θ values. Our adaptation method uses gradient-based optimization to update parameters θ over time, to be more consistent with the human's actions and intentions. This system allows users to teach robots to follow a new terrain boundary simply by driving along it. The robot's task performance also improves incrementally whenever the human corrects its mistakes.

We carried out boundary patrol experiments to compare our interactive adaptive system against both a non-adaptive *expert-tuned* variant, and plain teleoperation. In each experiment run, we collected various efficiency metrics, including automation failure rates, task durations, deviations from designated paths, manual intervention rates, and satisfaction ratings. Aggregated rankings of the 3 team configurations were then computed using the Kemeny-Young method.

In the first experiment, 15 users partnered up with a simulated aerial robot to track a straight highway stretch, a smooth forest path, and a curvy coastline segment. Aggregated rankings showed that our interactive adaptive robot performed comparably to the tediously expert-tuned setup, and both significantly outperformed plain teleoperation. In particular, the adaptive variant yielded greater efficiency during the coastline section, where the unpredictable curves in the terrain necessitated frequent updates to the tracker's parameters. In the second experiment, 7 users collaborated with a wheeled robot to patrol along footpaths, grass sides, and curbs on a university campus. The resulting aggregated efficiency rankings showed a consistent ordering, with users preferring our adaptive system, then the expert-tuned variant, and then plain teleoperation. These empirical results quantitatively substantiate the ability of our interactive adaptation strategy for increasing team efficiency, without requiring expert knowledge from the human collaborator.

4. TOWARDS TRUST-SEEKING ROBOTS

We are currently working to integrate our dynamic trust model into our online adaptation strategy, towards the end-goal of trust-aware adaptive robots that can actively seek out greater trust and efficiency. Specifically, during manual control, the amount of trust lost can be tied to the learning rate for our gradient-based adaptation method. This enables the robot to adapt more aggressively in response to severe cases of trust loss. In addition, when trust is low during autonomous control, the robot can initiate preemptive parameter optimization, to search for better settings in dire distrusting states. We will demonstrate this general methodology through multiple instances of trust-seeking adaptive robots. We expect our research to avert distrust and its fatal consequences [6], and enable high-efficiency and long-lasting collaborations for future human-robot teams.

5. REFERENCES

[1] B. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.

[2] M. Deisenroth, G. Neumann, and J. Peters. A survey on policy search for robotics. *Foundations and Trends in Robotics*, 2:1–142, 2013.

[3] M. Desai. *Modeling Trust to Improve Human-Robot Interaction*. PhD thesis, Computer Science Department, University of Massachusetts Lowell, 2012.

[4] J. Lee and N. Moray. Trust, control strategies and allocation of function in human-machine systems. *Ergonomics*, 35(10):1243–1270, 1992.

[5] J. Lee and K. See. Trust in automation: Designing for appropriate reliance. *Human Factors*, 46:50–80, 2004.

[6] R. Parasuraman and V. Riley. Humans and automation: Use, misuse, disuse, abuse. *H Factors*, 1:230–253, 1997.

[7] A. Xu and G. Dudek. Towards modeling real-time trust in asymmetric human-robot collaborations. In *Int. Sym. on Robotics Research (ISRR)*, 2013.

[8] A. Xu and G. Dudek. OPTIMo: Online probabilistic trust inference model for asymmetric human-robot collaborations. In *ACM/IEEE Int. Conf. on HRI*, 2015.

[9] A. Xu, A. Kalmbach, and G. Dudek. Adaptive parameter exploration (APEX): Adaptation of robot autonomy from human participation. In *Proc. of the IEEE Int. Conf. on Rob. and Auto. (ICRA)*, 2014.