

Human-Robot Cross-Training: Computational Formulation, Modeling and Evaluation of a Human Team Training Strategy

Stefanos Nikolaidis

Department of Aeronautics and Astronautics
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: snikol@mit.edu

Julie Shah

Department of Aeronautics and Astronautics
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: julie_a_shah@csail.mit.edu

Abstract—We designed and evaluated human-robot cross-training, a strategy widely used and validated for effective human team training. Cross-training is an interactive planning method in which a human and a robot iteratively switch roles to learn a shared plan for a collaborative task. Human subject experiments ($n = 36$) showed that cross-training provides statistically significant improvements in quantitative team performance measures, compared to standard reinforcement learning techniques. Additionally, significant differences emerged in the perceived robot performance and human trust. In this paper we briefly present our motivation and findings, which support the hypothesis that effective and fluent human-robot teaming may be best achieved by modeling effective practices for human teamwork.

I. INTRODUCTION

When humans work in teams, it is crucial for the team members to develop fluent team behavior. We believe that the same holds for robotics teammates, if they are to perform in a similarly fluent manner as members of a human-robot team. Learning from demonstration [1] is one technique for robot training that has received significant attention. In this approach, the human explicitly teaches the robot a skill or specific task. However, the focus is on one-way skill transfer from a human to a robot, rather than a mutual adaptation process for learning fluency in joint-action. In many other works, the human interacts with the robot by providing high-level feedback or guidance [4], but this kind of interaction does not resemble the teamwork processes naturally observed when human teams train together on interdependent tasks.

Our research leverages methods from human factors engineering, with the goal of achieving convergent team behavior during training and team fluency at task execution, as it is perceived by the human partner and is assessed by quantitative team performance metrics.

II. APPROACH

In our previous work [5], [6], we have computationally encoded the concept of shared mental models in the form of a Markov Decision Process (MDP), that captures the knowledge about the role of the robot and the human for a specific

task. We then proposed the entropy rate of the MDP as an objective metric to evaluate the convergence of the robot’s computational teaming model and the human mental model. Additionally, we designed a quantitative method to elicit the similarity between the mental model of human and robot, based on prior work on shared mental model elicitation for human teams.

Expert knowledge about the task execution is encoded in the assignment of rewards of the MDP, and in the priors on the transition probabilities that encode the expected human behavior. This knowledge can be derived from task specifications or from observation of expert human teams. However, rewards and transition probabilities finely tuned to one human worker are not likely to generalize to another human worker, since each worker develops his or her own highly individualized method for performing manual tasks. In fact, it has been shown in previous research that human teams whose members have similar mental models perform better than teams with more accurate but less similar mental models. Even if the mental model learned by observation of a team of human experts is accurate, the robot needs to adapt this model when asked to work with a new human partner. The goal then becomes for the newly formed human-robot team to develop a shared-mental model. One validated and widely used mechanism for conveying shared mental modes in human teams is “cross-training.” In [6], we emulated the cross-training process among human team-members by having the human and robot train together at a virtual environment.

III. HUMAN-ROBOT TEAMING EXPERIMENTS

We applied the proposed framework to train a team of one human and one robot to perform a place-and-drill task, as a proof of concept. The human’s role was to place screws in one of three available positions. The robot’s role was to drill each screw. This task is simple, but there is a sufficient variety on how to accomplish it among different persons. For example, some participants preferred to place all screws on a sequence from right-to-left and then have them drilled at the

same sequence, while others preferred to place and drill each screw before moving on to the next. The participants consisted of 36 subjects. Videos of the experiment can be found at: <http://tinyurl.com/9prt3hb>.

Each participant then did a training session in the ABB RobotStudio virtual environment with an industrial robot which we call “Abbie” (Figure 1). The participants were randomly assigned to two groups, Group A and Group B. Participants of Group A iteratively switched positions with the virtual robot, placing the screws at the forward phase and drilling at the rotation phase. Participants of Group B trained with the robot with the standard reinforcement learning approach, where the participant places screws and the robot drills at all iterations, with the participant assigning a positive, zero, or negative reward after each robot action [2].

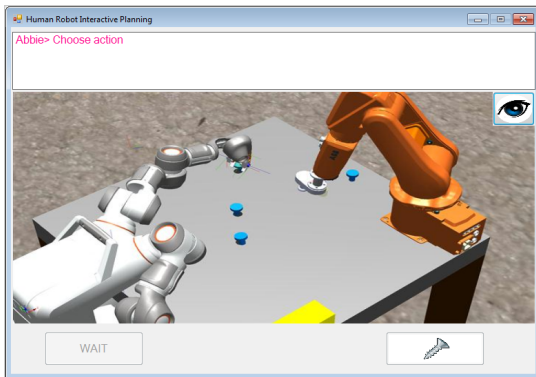


Fig. 1. Human-Robot Cross-Training using ABB RobotStudio Virtual Environment.

We then asked all participants to perform the place-and-drill task with the actual robot, Abbie (Figure 2). To recognize the actions of the human we used a Phasespace motion capture system of eight cameras, which tracked the motion of a Phasespace glove wore by the participant. Abbie executed the policy learned from the training sessions. The task execution was videotaped and later analyzed for team fluency metrics. Finally, all participants were asked to answer a post-experimental survey.




Fig. 2. Human-Robot Task Execution

IV. RESULTS

Results of the experiment showed that the proposed cross-training method outperforms standard reinforcement learning in quantitative measures of human-robot mental model convergence ($p = 0.04$) and mental model similarity ($p = 0.03$). Additionally, the post-experimental survey showed that participants of Group A agreed more strongly that Abbie learned their preferences, compared to participants of Group B, and trusted Abbie more ($p < 0.01$), in accordance with prior work [7]. We also elicited the fluency of the teamwork by measuring the concurrent motion of the human and robot and the human idle time [3] during the task execution phase that succeeded the human-robot training process. We observed an increase of 71% in concurrent motion ($p = 0.02$) and a decrease of 41% in human idle time ($p = 0.04$). One possible explanation for this difference is that cross-training engendered more trust in the robot, and thereby participants of Group A had more confidence to act while the robot was moving. In some cases, the increase in idle time was caused because the participant was not sure on what the robot would do next, and therefore waited to see. In other cases, the robot had not learned correctly the human preference and did not act accordingly, with the result of forcing the human to wait, or confusing the human team-member.

V. CONCLUSION

We presented our motivation on using methods from prior studies on human teamwork, and on applying them on a team of a human and a robot. Recent results provide the first evidence that human-robot teamwork is improved when a human and robot train together by switching roles, similarly to practices observed in human teams. In [6], we focused on a simple place-and-drill task, as a proof of concept. We are currently planning on extending the framework on more complex tasks, and on using the robot uncertainty about the human’s next action to influence the motion planning parameters for a robot working alongside a person. We are also extending the computational formulation of the robot’s teaming model to a POMDP framework that incorporates information-seeking behavior.

REFERENCES

- [1] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, “A survey of robot learning from demonstration,” *Robot. Auton. Syst.*, vol. 57, no. 5, pp. 469–483, May 2009.
- [2] F. Doshi and N. Roy, “Efficient model learning for dialog management,” in *Proc. HRI*, Washington, DC, March 2007.
- [3] G. Hoffman and C. Breazeal, “Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team,” in *Proc. HRI*. New York, NY, USA: ACM, 2007, pp. 1–8.
- [4] W. B. Knox and P. Stone, “Reinforcement learning from simultaneous human and mdp reward,” in *Proc. AAMAS*, June 2012.
- [5] S. Nikolaidis and J. Shah, “Human-robot interactive planning using cross-training: A human team training approach,” in *Proc. Infotech@Aerospace*, 2012.
- [6] —, “Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy,” in *Proc. HRI*, 2013.
- [7] J. Shah, J. Wiken, B. Williams, and C. Breazeal, “Improved human-robot team performance using chaski, a human-inspired plan execution system,” in *Proc. HRI*. New York, NY, USA: ACM, 2011, pp. 29–36.