Developing Effective Robot Teammates for Human-Robot Collaboration

Bradley Hayes and Brian Scassellati
Yale University
Department of Computer Science
New Haven, CT 06511

Introduction

Developing collaborative robots that can productively operate out of isolation and work safely in uninstrumented, human-populated environments is critically important for advancing the field of robotics. The development of such systems, those that handle the dynamics of human environments and the complexities of human interaction, is a strong focus within Human-Robot Interaction and involves underlying research questions deeply relevant to the Artificial Intelligence community.

Especially in domains where modern robots are ineffective, we wish to leverage human-robot teaming to improve the efficiency, ability, and safety of human workers. As a community, we desire to create collaborative robots that can provide assistance when useful, remove dull or undesirable responsibilities when possible, and provide instruction or guidance when necessary.

Doing so requires addressing a multitude of deep and challenging research questions. Inferring the intentions of one's collaborators is critically important for effective teaming, yet remains an extremely complex problem even under instrumented environments. Optimally planning one's own actions under uncertainty is also a necessary field of study, as human populated environments are dynamic and often contain nigh-unpredictable actors. Task comprehension and knowledge transparency, particularly building shared mental models with teammates, is crucial to multi-agent collaboration. Finally, user modeling plays a large role in optimizing team behavior, yet remains an open problem as feature selection, extraction, analysis, and exploitation are difficult to generalize across scenarios.

Research Overview

Our work focuses on creating agents capable of human-robot teamwork, in particular the case where a team of humans and robots are working together towards a common goal. This research exists at the intersection of a host of important robotics problems (Hayes and Scassellati 2013a). These include learning motor primitives from demonstration (LfD), learning hierarchical task networks (Hayes and Scassellati 2014a), performing multi-agent planning and state estimation, inferring other agents' intentions (Hayes and Scassellati 2013b), collaborative manipulation, and legibly conveying internal knowledge and understanding. In particular, we focus on collaboration between a lead worker and robotic assistant, complementing prior work that develops collaborative robots as peers (Knepper et al. 2013) and those that learn from demonstration (Konidaris et al. 2011; Cakmak and Thomaz 2012).

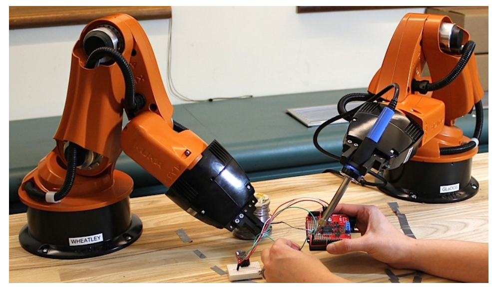

Figure 1: One human-robot teaming domain we focus on is the Collaborative Workbench platform, designed for shared workspace, close proximity human-robot teaming exercises.

The research we perform is singularly focused on building the technology required to produce a capable robot assistant that comprehends complex tasks and is trainable by non-technical subject matter experts. To be effective, this assistant must be able to learn to anticipate the physical and materials-related needs of its collaborators, adapting to and constantly evaluating its own desired plans against the preferences of its teammates. Further, our work introduces methods for an assistive robot teammate to become an instructor, using learned task information to autonomously generate strategies for training novice teammates.

Human-Robot Collaboration

Two popular paradigms for collaboration are 'leader-follower' and 'equal partners'. In leader-follower teaming, one agent drives the progression of the task while the follower facilitates subtask completion. In equal partners, each agent takes a lead role in task execution. Both paradigms require an awareness of one's collaborators and the ability to appropriately react to them. Our work focuses primarily on the leader-follower paradigm of human-robot teamwork. Within this model of collaboration, we examine two distinct roles for a robot to assume: assistant and mentor. As each teaming strategy is dependent upon the robot having a thorough understanding of the task being performed, task comprehension is a core component of our work.

Task Comprehension

By leveraging a combination of kinesthetically trained dynamic movement primitives (DMPs), Partially Observable Markov Decision Processes (POMDPs), and Active Learning, we develop a robot capable of learning hierarchical task networks (HTNs). While collaborating, maintaining situational awareness and accurately predicting coworker intent is critical. Thus, using a novel graph transformation algorithm, we are able to use these HTNs to build hierarchical, goal-centric (as opposed to environment-state centric) POMDPs that are used to infer co-workers' intentions.

Having developed the ability to transform a task into a hierarchical network of subtask-level goals, we have the information necessary to design algorithms for enabling productive assistive and instructive behaviors.

Assistive Roles

We predominantly consider application domains in which a human is performing a task while physically sharing a workspace with a robot (Figure 1). The central contribution of this research is a means of learning different types of assistive behaviors and the contexts in which to apply them (Hayes and Scassellati 2014b). Learned behaviors take the form of DMPs, which can range from simple materials stabilization constraints to joint object manipulations. Primary concerns for this operating mode include building shared mental models with teammates (having comparable conceptions of the task steps and progression) and performing social modeling to learn user preferences (e.g., proxemics or subtask completion ordering).

Our work takes these subproblems into account while developing algorithms that allow for the robot to associate LfD-acquired assistive behaviors with relevant nodes in a HTN. Collaborators are able to teach a robot, via gesture-based instruction or kinesthetic manipulation, how to be helpful throughout task execution. These assistive behaviors can be generalized by applying them across similar substructures within the HTN. We utilize human-in-the-loop reinforcement learning to build personalized user models for determining which assistive behaviors to apply at a given HTN node, tracking features such as 'duration to action response' and 'user {accepted|rejected} action'. This work contributes towards enabling an experienced worker to improve her efficiency, work quality, and personal safety throughout task completion.

Instructive Roles

A distinct but related challenge we address is developing algorithms that allow a robot to guide a teammate through a task, providing training for required skills and action sequences. In these scenarios, we imagine introducing a novice or uninformed agent to a novel task. Using information from the robot's HTN and the learned assistive behaviors, the robot is able to provide materials and usage context to its teammates. Of fundamental importance within this work is the ability to generate instructional behaviors by autonomously adapting existing functional, task execution-based information.

When providing instruction, It is important to use appropriate levels of abstraction to guide the learner at an acceptable pace. Merely providing instruction that specifies each action with precise detail will quickly frustrate and bore the learner, losing engagement. Overzealous abstraction leads to confusion and potentially invalid task executions. To appropriately balance this abstraction, we use timing information and observed object manipulations to evaluate learner progress and subtask competency. With this information, we base the sophistication and explicitness of instruction upon hypotheses of the learner's level of understanding.

Acknowledgments

This work is supported by Office of Naval Research grant #N00014-12-1-0822.

References

Cakmak, M., and Thomaz, A. 2012. Designing robot learners that ask good questions. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, 17–24. ACM.

Hayes, B., and Scassellati, B. 2013a. Challenges in shared-environment human-robot collaboration. In *Proceedings of the Collaborative Manipulation workshop at the 8th ACM/IEEE International Conference on Human-Robot Interaction*. IEEE.

Hayes, B., and Scassellati, B. 2013b. Improving implicit communication in mixed human-robot teams with social force detection. In *2013 IEEE Third Joint International Conference on Development and Learning and Epigenetic Robotics*. IEEE.

Hayes, B., and Scassellati, B. 2014a. Discovering task constraints through observation and active learning. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*.

Hayes, B., and Scassellati, B. 2014b. Online development of assistive robot behaviors for collaborative manipulation and human-robot teamwork. In *Proceedings of the "Machine Learning for Interactive Systems" (MLIS) Workshop at AAAI 2014*.

Knepper, R. A.; Layton, T.; Romanishin, J.; and Rus, D. 2013. Ikeabot: An autonomous multi-robot coordinated furniture assembly system. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, 855–862. IEEE.

Konidaris, G.; Kuindersma, S.; Grupen, R. A.; and Barto, A. G. 2011. Autonomous skill acquisition on a mobile manipulator. In *Twenty-Fifth Conference on Artificial Intelligence*, 1468–1473.