Social Hierarchical Learning

Bradley Hayes
Department of Computer Science
Yale University
bradley.h.hayes@yale.edu

Summary

My dissertation research focuses on the application of hierarchical learning and heuristics based on social signals to solve challenges inherent to enabling human-robot collaboration. I approach this problem through advancing the state of the art in building hierarchical task representations, multi-agent task-level planning, and learning assistive behaviors from demonstration.

Introduction

The cages and physical barriers that once isolated robots from contact with humans are being replaced with sensing technology and algorithms. As such, collaborative robotics is a fast-growing field of research spanning many important real-world robotics and artificial intelligence challenges. These include learning motor skills from demonstration, learning hierarchical task models, multi-agent planning under uncertainty, and intention recognition. My dissertation research focuses on how to develop and apply hierarchical learning methods that leverage social signals as heuristics to make collaboration between humans and robots safe, productive, and efficient. In particular, my work focuses on methods applicable to sequential motor tasks.

Other researchers that work within this domain have focused on effective multi-robot team planning (Dogar et al. 2014), efficiently solving single-agent sequential motor task planning problems (Tomas Lozano-Perez 2014), and investigating theoretical properties of multi-agent problem representations (Dibangoye et al. 2014). Prior related work concentrating on human-robot interactions has focused on building shared mental models amongst collaborators in mixed human-robot teams (Nikolaidis and Shah 2013) and increasing task knowledge from social interactions (Kollar et al. 2013).

My research can be split into three inter-related components: learning task structure, performing socially aware multi-agent task planning, and learning supportive behaviors for human-robot collaboration.

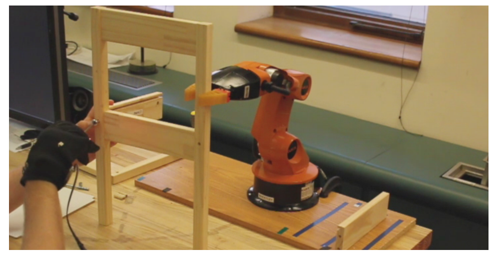

Figure 1: Collaborative execution of a furniture construction task with a mixed human-robot team. In this situation, the team is interacting in a “leader-assistant” paradigm, with the robot performing supportive actions to increase the efficiency of the human worker.

Learning Task Structure

Achieving collaboration between humans and robots is contingent upon providing adequate solutions to many challenging problems (Hayes and Scassellati 2013), including state estimation, goal inference, and action planning. Scalable, efficient tools that offer solutions to many of these problems are available, but rely on the prior availability of a rich, hierarchical task model. While some techniques exist for building these models, they place heavy requirements on the available information encoded in the task representation.

In many cases, especially those where the predominant mode of task specification is user demonstration, the extraction and formalization of this information is exceptionally difficult and can easily be intractible given the available sensor data. Accordingly, one area of my research concerns developing robots that engage in active learning, participating in the learning process by guiding their instructors to provide more informative examples to extend structural task knowledge (Hayes and Scassellati 2014a).

In work under review, I present an algorithm for constructing flexible hierarchical task networks, derived from sub-goal ordering constraints within task networks (e.g., Semi-Markov Decision Processes), that are compatible with demonstration-based skill acquisition techniques. This compatibility emerges from the minimal symbolic requirements and insights levied upon the chosen goal and/or motor primitive representation. The resulting hierarchical task structure has many practical applications within human-robot collaboration. In this work, I present a proof-of-concept goal inference solution, using hierarchical Hidden Markov Models directly constructed from the generated task hierarchy, achieving multi-resolution goal inference improving state estimation accuracy while maintaining the computational benefits expected of hierarchical approaches. This multi-level abstraction provides the opportunity for cross-task knowledge transfer, allowing an agent to reuse higher level policies from prior experiences in future scenarios.

Performing Collaborative Task-level Planning

Given a hierarchical task model, multi-agent, complex task planning can be transformed into a problem of role assignment. In work that is in progress, I integrate the interpretation of human social signals into a multi-agent planner. This is expected to improve the fluency of collaboration, as people utilize a host of non-verbal social signals to communicate intent, needs, and preferences during human-human collaboration that are typically ignored during robot interactions.

Consequently, we take an approach congruous with (Tomas Lozano-Perez 2014), framing the sequential motor planning task as a coupled symbolic planning problem and geometric constraint satisfaction problem. Our work makes three novel contributions: we generalize this constraint-based formulation to dynamic environments, incorporate social signals into the planner to influence goal prioritization and sub-task assignment, and implement a hierarchical constraint approach to improve performance towards achieving feasibility in live contexts.

Future work that I have planned includes extending the fidelity and depth of social signal understanding to natural language processing. By using the autonomously constructed hierarchical task model and live execution contexts as a grounding for language, verbal communication can be leveraged as a reliable, real-time planning cue.

This work is extremely relevant to the future of collaborative robotics platforms, as taking steps towards solving the multi-agent planning problem at both the symbolic and geometric level is paramount for effective human-robot teams in dynamic environments.

Learning Supportive Behaviors

The final component of my dissertation research involves the acquisition of supportive motor primitive actions from task-level planners or collaborator demonstration. Such actions can be used in an assistive capacity throughout a team interaction (Hayes and Scassellati 2014b), following the “leader-assistant” collaboration paradigm. In particular, this work seeks to develop algorithms enabling a capable robot assistant to learn and utilize supportive motor primitives in a live task execution context, reducing the cognitive load or dexterity requirements of a co-worker’s current task. These supportive motor primitives can be classified as belonging to five distinct categories: materials stabilization, materials retrieval, collaborative object manipulation, awareness improvement, and task progression guidance. Each requires a different basis of approach that can be shaped and specialized with instructor-provided demonstrations.

More formally, this research involves policy learning from demonstration, where the agent is attempting to derive an execution policy specific to a presented task context (subgoal). This can be conceptualized as learning a set of complementary motor primitives that may be associated with arbitrary options from a Semi-Markov Decision Process. In practice, these motor primitives are encoded in a Partially Observable Markov Decision Process associated with a state (representing a task sub-goal) within a Hidden Markov Model used to track task progress. An example of this (Figure 1) would be executing an assistive behavior (with materials stabilization action) that is associated with the “attach left support” subgoal of a chair construction task.

This has immediate applicability to an array of task domains, particularly those where a robot is incapable by itself or where it would be unsafe or inefficient for a lone human.

References

Dibangoye, J. S.; Amato, C.; Buffet, O.; and Charpillet, F. 2014. Exploiting separability in multiagent planning with continuous-state mdps. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 1281–1288. International Foundation for Autonomous Agents and Multiagent Systems.

Dogar, M.; Knepper, R. A.; Spielberg, A.; Choi, C.; Christensen, H. I.; and Rus, D. 2014. Towards coordinated precision assembly with robot teams. In *Proceedings of the International Symposium of Experimental Robotics*.

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

Hayes, B., and Scassellati, B. 2014a. Discovering task constraints through observation and active learning. In *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*.

Kollar, T.; Perera, V.; Nardi, D.; and Veloso, M. 2013. Learning environmental knowledge from task-based human-robot dialog. In *IEEE International Conference on Robotics and Automation*, 4304–4309. IEEE.

Nikolaidis, S., and Shah, J. 2013. Human-robot cross-training: computational formulation, modeling and evaluation of a human team training strategy. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot Interaction*, 33–40. IEEE Press.

Tomas Lozano-Perez, L. P. K. 2014. A constarint-based method for solving sequential manipulation planning problems. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*.