Social Hierarchical Learning

Bradley Hayes
Department of Computer Science
Yale University

***Abstract*—I present the outline for my dissertation work, Social Hierarchical Learning (SHL). SHL leverages and extends the capabilities of state-of-the-art hierarchical learning (HL) systems to operate under realistic human-robot interaction domains. SHL is designed for multi-agent coordination with humans in the loop. Traditional HL systems excel when presented tasks in which the agent has full environmental awareness and control, yet such systems are not intended to handle elements outside the agent's control interacting within its workspace. SHL is designed to provide the required flexibility in task decomposition and assignment for successful human-robot collaboration.**

I. Introduction

As recent years have demonstrated, the introduction of robot systems to collaborative environments requires a substantial engineering effort aimed at not only producing capable and robust robots but also requiring careful and planned interaction methodology for human operators to utilize robotic resources. In most cases, the dominant metaphor for the operation of these robots is as a co-worker or teammate, operating side-by-side with human personnel, rather than as a fully autonomous system operating in isolation [1]. To construct flexible systems that can adapt to the changing needs of daily operations, researchers have proposed to construct robots that learn basic skills that can be re-used in multiple tasks and under varying conditions. To this end, hierarchical learning is designed to support skill abstraction and to provide a level of portability to acquired knowledge that improves performance on distinct future tasks.

While traditional hierarchical learning has been successfully applied to a number of real-world robotic tasks in which the robot acts autonomously in isolation, this method suffers a critical weakness when applied to collaborative scenarios. Division of responsibilities, role identification, and joint actions are unaddressed, leaving most hierarchical systems inapplicable in collaborative domains. Where traditional hierarchical learning assumes an isolated, autonomous robot that learns and performs on its own, Social Hierarchical Learning enables and develops a collaborative robot that pairs social modeling with portable skills from human guidance and engages in tasks with human co-workers.

II. Background and Motivation

Hierarchical learning provides a variety of approaches permitting traditional reinforcement learning algorithms to handle complex or large problem spaces. This is accomplished by introducing levels of abstraction, encapsulating sequences of primitive actions or sequences of already-encapsulated actions into higher-level sequences. Hierarchical reinforcement learning seeks to develop methods by which an agent can autonomously acquire its own higher-level skills [2], [3]. Recent HL systems have begun to incorporate skill portability. This enables an agent to develop a library of skills, commonly through exploration, and determine which are problem specific and which can be transferred beyond the current domain [4].

While there have been substantial efforts devoted to the development of HL algorithms, these systems are not well suited to collaborative human-robot interaction (HRI) tasks. HL algorithms are typically designed for robots that have complete control over their selection of actions, and have been extended to multi-robot cooperative domains [5]. In HRI scenarios, robots must continuously adapt to the preferences and needs of the human partner while at the same time accounting for the limits of the robots own capabilities. Further, the robot may be capable of executing only parts of the joint task to be accomplished. Finally, even for tasks that the robot could perform, the actions of the human partner will influence which parts of the task the robot should engage in, and in which order those tasks should be attempted. These cases are not typically designed for in traditional HL systems.

III. Social Hierarchical Learning

SHL accomplishes hierarchical learning for socially cooperative tasks between one or more robots and one or more humans operating in the same physical space on the same tasks. Reinforcement learning and learning by demonstration are leveraged to acquire basic skill competency. Second, the nature of the overall task presented to the system is learned, along with a decomposition of subtasks (a ”plan”). Finally, the system learns how best to assign roles in real-time, adapting SHL agents to collaborate with human co-workers to improve efficiency and performance while executing a plan.

The SHL test environment is a workbench with two KUKA youBot arms mounted to it (Fig. 1). The system is available as a ROS stack, utilizing existing libraries for collision detection and scene modeling. Current SHL work centers on learning effective skill execution policies that are compatible with a human operating in the same physical space. Thus, given an existing comprehension of how to execute a skill, this work determines the best way to adapt existing knowledge for safe execution in an environment with humans by adding heuristics derived from socially predictive models for potential external agents.

Social Hierarchical Learning consists of three phases: primitives acquisition, plan decomposition, and cooperative execution.

A. *Primitives Acquisition*

In the first phase of SHL, human-guided reinforcement learning and learning from demonstration is used to acquire component sub-skills. This allows non-experts to become effective instructors, simultaneously leveraging humans foresight while remaining independent of complex natural language processing requirements. A human operator interacting with an SHL-enabled agent must first teach the robot the basic skills that it lacks. Due to the portability of skills from task to task, it is reasonable to suppose that the robot may already know how to perform some relevant actions, but not others. After a small number of demonstrations, the robot should be able to reproduce a similar action in similar settings, but may not have an entirely accurate representation of the concept when taken to new environments or settings. Mistakes can be negatively reinforced by feedback or repeated demonstrations from the human partner, eventually eliminating undesirable behavior and resulting in proficiency. Choosing not to train the robot on such primitives limits the eligible roles for it in the final cooperative execution phase.

B. Plan Decomposition

The second phase involves learning the structure of the task to be solved. Utilizing methods akin to learning from demonstration, the system learns to sequence primitive actions to achieve a complex hierarchical task representation. The primary focus of this research is discovering, manipulating, and optimizing the collaborative structure of the task, rather than goal state discovery or primitive action learning. Research within robotics as it pertains to learning from demonstration typically involves a robot mimicking an action undertaken by a single human or robot. Modeling an entire interaction, inclusive of all actors from initiation to goal state, is helpful but not altogether necessary for a SHL system to succeed. A robot placed in a situation where it was fully trained on the entire interaction would determine its role quicker and react more effectively to those it is working with. The same robot placed in a situation on which it is not fully trained could accomplish this learning through its own experience by observing the humans it is working with.

This process of observation results in a skill tree and sample successful traversal (a sample plan ”solution”), with social metadata indicating potential roles that may be assigned within a particular skill tree traversal. This social metadata adds context to known skills and helps restrict the action search space. A substantial component of SHL-based research involves establishing a suitable representation for this skill tree that contains all of the required social modeling and metadata for constructive collaboration. The hierarchical structure of the task is learned simultaneously with the relevant social metadata.

C. Cooperative Task Execution

The final phase of SHL involves learning to produce live role assignments in an effective and safe way. Reinforcement learning is applied to teach the system how to divide labor amongst autonomous workers in response to the human workers' actions in a socially optimal way. This feedback regarding the system's assignment and parallelization of tasks to agents guides the autonomous agents' resource allocation. The primary challenge of this phase is analyzing the social metadata and applying it to both the skill tree and proposed traversal of it (plan solution). The result of this process is a role assignment tree, a flexible assignment structure that can compensate for unknowns and uncontrollable agents in the overall task completion plan.

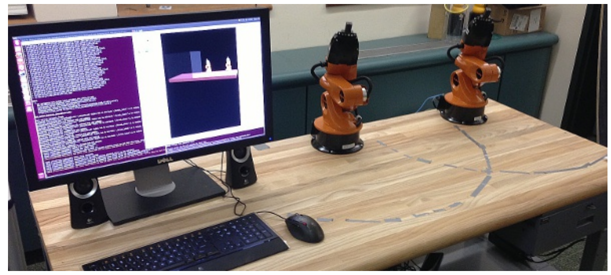

Fig. 1. Social Hierarchical Learning Workbench

IV. Impact

Flexible, dynamic cooperation on this level is novel and would greatly benefit hierarchical learning systems and planners. Real-time, reactive role determination and learned task execution combine to form a powerful system enabling side-by-side collaboration between humans and robots. Roles may inherently carry restrictions within this role assignment paradigm, as it can be learned that certain tasks may be better suited for robots to complete and that some tasks should exclusively be performed by humans.

The primary contribution of the proposed work is to take steps toward more natural, task-centered, shared space, peer-to-peer human-robot interactions. The SHL architecture allows for a cognitively compatible representation of task composition, which is essential for both allowing transparent interactions and for constructing task allocations dynamically.

A secondary contribution of this work is an architecture for collaborative work. This architecture allows for a representation of structured tasks that can be manipulated for variations in individual preferences or capabilities. Understanding the structure of a task and particular variations, may also lead to a better understanding of cognitive load (and especially overload) in demanding task scenarios. Future experiments based on comparisons of task load could lead to a more general model of how robots can collaborate efficiently with humans, reducing the time required for plan execution and cognitive load requirements placed upon human co-workers.

References

[1] J. Shah, J. Wiken, B. Williams, and C. Breazeal, “Improved human-robot team performance using chaski, a human-inspired plan execution system,” in *Proceedings of the 6th international conference on Human-robot interaction*. ACM, 2011, pp. 29–36.

[2] A. Barto and S. Mahadevan, “Recent advances in hierarchical reinforcement learning,” *Discrete Event Dynamic Systems*, vol. 13, no. 4, pp. 341–379, 2003.

[3] A. Stout, G. Konidaris, and A. Barto, “Intrinsically motivated reinforcement learning: A promising framework for developmental robot learning,” 2005.

[4] G. Konidaris, S. Kuindersma, R. Grupen, and A. Barto, “Autonomous skill acquisition on a mobile manipulator,” in *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence*, 2011, pp. 1468–1473.

[5] M. Ghavamzadeh and S. Mahadevan, “Learning to communicate and act using hierarchical reinforcement learning,” in *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 3*. IEEE Computer Society, 2004, pp. 1114–1121.