Apprenticeship Scheduling for Human-Robot Teams in Manufacturing

Matthew C. Gombolay
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, Massachusetts 02139
gombolay@csail.mit.edu

Introduction

Traditional uses of robotic technology in manufacturing have seen robotic work physically caged-off and separated from human workers. Automotive manufacturing is a success case for manufacturing automation. However, the human element is still an incredibly important facet of automotive manufacturing. Manual work still counts for 50% of the build process and requires two-thirds of the factory footprint. The process to build a car remains a manual task, which requires the dexterous skills of people. Recent advances in robotic technology are opening up the possibility of integrating mobile robotic assistants into areas of the factory that have traditionally been reserved for human workers. Research in localization and mapping, computer vision, control systems, and learning-from-demonstration is allowing these robots to support non-value added tasks (e.g., fetching of parts) and even value-added tasks in some cases (e.g., welding, drilling, etc.).

Successfully introducing this advanced robotic technology to work on human-robot teams requires an algorithm that can both safely choreograph the activity of these teams in real-time and generate robotic behavior that gains the trust of the human workers and managers. Human workers often develop a sense of identity, security, and purpose from their jobs or roles in factories. Similarly, human workers develop mental models or expectations of how the team as a whole should work together. These patterns and expectation of workflow can change from team to team even when working on the same job. Thus, giving robots the ability to understand how to complete the set of tasks in a way that is intuitive to their human teammates is paramount to the success of integrating robots into human workspaces.

Developing an user-interface by which humans can command or directly operate robotic assistants also presents a challenge to the success of integrating robots into the human workspace. The human-robot interface has long been seen as a major bottleneck for the success of these robotic systems (Parasuraman, Mustapha, and Hilburn 1999), even though much effort has been conducted to improve the fluidity of the interaction (Clare et al. 2012). Other researchers have taken the tack of explicitly formulating the problem with mathematical constraints and optimization criteria (Ryan et al. 2013). This approach is especially important because, in many domains, it is often difficult or impractical to hand-code the appropriate constraints and rules-of-thumb that human experts learn over years of experience in the field. For example, a study of aircraft carrier flight deck operations showed that the heuristics veteran operators use outperform mathematical optimization (Ryan et al. 2013).

Apprenticeship Scheduling

We as roboticists need the ability to efficiently learn the how to coordinate a fluent team from the human experts who have gained knowledge and experience over years at their tasks. My PhD dissertation will be developing the right mathematical models and computational techniques to autonomously learn the implicit heuristics of expert operators to coordinate resources in real time. The first challenge is developing fast methods for the dynamic coordination of human-robot teams that can adapt robot behavior in real-time to support fluent teamwork. Second, we need to understand how people change their behavior based on parameters such as their level of control over team coordination and the composition of the team (e.g., human-human versus human-robot) to help discover the correct mathematical representations to model this behavior. The third challenge is developing computational techniques to learn how people schedule team activities, encode that knowledge in robotic systems, and validate that the learned behavior improves team fluency.

Fast Methods for Human-Robot Team Coordination

The first step in developing an apprentice scheduler for human-robot teams is creating scalable computational models and techniques to reason one the space of possible schedules. Problems of interest in manufacturing consist of upper and lowerbound temporal constraints, spatial constraints, and heterogeneous agent capabilities. Approaches in prior work often decompose the problem to gain computational tractability by having each agent schedule their own tasks. However, when agents work in close proximity with shared spatial resources, each commitment an agent makes directly affects the utility of the rest of the team. Methods that reason about each agent separately thus lose their ability to efficiently reason when solving problems of interest. I developed a computational technique that efficiently reasons about each agents use of shared resources to provide a real-time coordination of robots working alongside human workers. An example scenario can be seen at `http://youtu.be/_qb2_jJID5c`. To ensure that agents do not violate temporal and spatial constraints, I formulated an analytical schedulability test that upperbounds the temporal and spatial resources required to complete subsets of the build process (Gombolay, Wilcox, and Shah 2013). This technique for real-time scheduling of human-robot teams serves as the first step towards developing the correct modeling for apprentice scheduling.

Human Factors of Decision-Making Authority and Team Composition

Before modeling the human decision-making process, I conducted a set of human-subject experiments to understand how people would react to working alongside and sharing decision-making authority with robotic teammates. In our experiment, a human subject worked on a human-robot team to complete a set of tasks analogous to a manufacturing environment. We presented three scenarios to the participant: 1) subjects could allocate work to each member of the human-robot team, 2) subjects could control only the tasks they would perform, and 3) the robotic agent would allocate the work to each member of the team. 3) I found that humans make qualitatively different scheduling decisions based on the level of control they have over their own work and the work of their robot teammates (Gombolay et al. 2014). I am currently conducting a set of experiments where the subject works on a human-human team to isolate the effects of working with a human versus a robot. The insights I have gained from these experiments will help inform the design of the right mathematical representation of human decision-making for coordinating human-robot teams.

Learning Implicit Constraints and Goals for Team Coordination

Learning the implicit rules-of-thumb, heuristics, hard and soft constraints of a complex team coordination task is a challenging problem. Many approaches to learning from observation assume that experts are behaving optimally. However, human factors studies of decision-making show that experts in time-critical domains often make decisions that satisfy relevant constraints rather than taking the time to search for the optimal solution. For example, studies of firefighters show that people in these domains find previous experiences that most resemble the current scenario and use that as the basis for their actions (Klein 1993). Furthermore, different experts can generate different but equally productive plans for managing resources in a factory.

One common approach to learning from observation that has been quite successful is Inverse Reinforcement Learning. Researchers have extended the capability of IRL algorithms to be able to learn from operators with differing skill levels (Ramachandran and Amir 2007) and identify operator subgoals (Michini and How 2012). IRL is able to leverage the structure of the Markov-Decision Process to bind the rationality of the agent. However, resource optimization or scheduling in manufacturing is highly non-Markovian: the next state of the environment is dependent upon the history of actions taken to arrive at the current state and the current time. Some researchers have tried to extend the traditional Markov Decision Process to characterize temporal phenomena, but these techniques do not scale to solving large-scale coordination of human-robot teams.

Attending AAAI, the premier conference in artificial intelligence, would help me to understand the state-of-the-art in artificial intelligence, mathematical modeling of human and machine decision-making, and robotics. Exchanging ideas and developing collaborations with researchers in the AAAI community would be an invaluable opportunity to help me realize the goals of my dissertation. By providing manufacturers and other practitioners of robotics with an apprentice scheduler that can learn how to coordinate manufacturing resources based on the expertise of current human operators, we can enable the successful integration of robots into the human workplace both in terms of human-robot team efficiency and the likelihood that the robotic system will be adopted and trusted by its human teammates.

References

Clare, A. S.; Cummings, M. M. L.; How, J. P.; Whitten, A. K.; and Toupet, O. 2012. Operator objective function guidance for a real-time unmanned vehicle scheduling algorithm. 9:161–173.

Gombolay, M. C.; Gutierrez, R. A.; Sturla, G. F.; and Shah, J. A. 2014. Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. In *Proceedings of the Robots: Science and Systems (RSS)*.

Gombolay, M. C.; Wilcox, R. J.; and Shah, J. A. 2013. Fast scheduling of multi-robot teams with temporospatial constrints. In *Proceedings of the Robots: Science and Systems (RSS)*.

Klein, G. A. 1993. *A recognition-primed decision (RPD) model of rapid decision making*. New York, NY, USA: Ablex Publishing Corporation.

Michini, B., and How, J. 2012. Bayesian nonparametric inverse reinforcement learning. In Flach, P.; De Bie, T.; and Cristianini, N., eds., *Machine Learning and Knowledge Discovery in Databases*, volume 7524 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. 148–163.

Parasuraman, R.; Mustapha, M.; and Hilburn, B. 1999. Adaptive aiding and adpative task allocation enchance human-machine systems. *Automation technology and human performance: Current research and trends* 119–123.

Ramachandran, D., and Amir, E. 2007. Bayesian inverse reinforcement learning. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence*, IJCAI'07, 2586–2591. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

Ryan, J. C.; Banerjee, A. G.; Cummings, M. L.; and Roy, N. 2013. Comparing the performance of expert user heuristics and an integer linear program in aircraft carrier deck operations. *IEEE Transaction on Cybernetics* PP(9).