Challenges in Shared-Environment Human-Robot Collaboration

Bradley Hayes
Department of Computer Science
Yale University
New Haven, CT 06511
bradley.h.hayes@yale.edu

Brian Scassellati
Department of Computer Science
Yale University
New Haven, CT 06511
scaz@cs.yale.edu

***Abstract*—We present and discuss four important yet underserved research questions critical to the future of shared-environment human-robot collaboration. We begin with a brief survey of research surrounding individual components required for a complete collaborative robot control system, discussing the current state of the art in Learning from Demonstration, active learning, adaptive planning systems, and intention recognition. We motivate the exploration of the presented research questions by relating them to existing work and representative use cases from the domains of construction and cooking.**

I. Introduction

Robots have the potential to revolutionize many domains, spanning all manner of industry from manufacturing to healthcare. Today's robots operate autonomously with speed and precision within tightly controlled, isolated environments. The next step in personal and industrial robotics is to move robots out of isolation and into collaborative relationships, operating side-by-side with human personnel. As demonstrated in recent years, the introduction of these systems to real-world environments requires a substantial engineering effort with a carefully planned interaction methodology, so that human operators can effectively utilize robotic resources. We include in this paper an exposition of what we believe to be some of the most important research questions yet to be fully addressed within shared-environment collaborative robotics. We highlight the challenges associated with these research questions and identify existing and future work in the field.

Collaborative operation is ripe for exploration and innovation, opening the possibility for widespread adoption of robots into problem domains for which they may currently be perceived as unsuitable or unready. Even with the current state-of-the-art in skill acquisition and execution, robots have difficulty performing at their full potential when removed from the typically well-controlled environment of the lab.

In this paper, we will use the term *collaborative task execution* to describe an agent autonomously performing a task either collaboratively with or in the presence of other agents, while respecting any associated social roles and divisions of responsibility. Within this definition, we focus exclusively on fully autonomous robots and humans working as a team, as opposed to teleoperated agents. A *skill* is defined as a temporally extended action similarly to options in reinforcement learning [1], and is assumed to minimally include a set of known preconditions, expected post-conditions, and known goal states. We define a *task* as a tree of skills and a *plan* as an arrangement of skills resulting from a complete traversal of that tree, respecting any ordering constraints. A plan constructed to include multiple agents may include parallel branches of skill execution, subject to the ordering constraints of the individual skills involved.

Requirements for robot control systems that are capable of engaging in collaborative behaviors with humans are incredibly complex. Such systems require many components to operate safely and properly, often leveraging the state-of-the-art in Learning from Demonstration (LfD), intention detection, speech recognition, non-verbal communication, planning systems, physical manipulation, and more. In this paper, we outline four major challenges in building complete autonomous, socially collaborative systems and discuss how the state-of-the-art from the fields of HRI, LfD, hierarchical learning, and reinforcement learning can be synthesized to help solve these problems. Attempting to build a socially collaborative system forces the designer to address issues of safety, user modeling, environment sensing, and various teamwork-related human factors in addition to the already numerous technical challenges inherent to robotics involving real-world manipulation [2].

II. Background and Related Work

Most existing work relevant to human-robot collaboration has focused on two topics: engaging non-technical users and skill acquisition. Developing capable systems that are accessible to non-roboticists is a primary concern for collaborative robotics. One research area essential for accessible collaborative systems that has received a great deal of attention is skill transfer. Shared-space human-robot collaboration can be a powerful enabler for skill transfer between humans and robots. In particular, collaborative tasks provide great opportunities for applications of learning from demonstration. LfD includes mechanisms for enabling skill transfer between humans and robots, producing systems that learn from skill executions led by humans or other agents [3]. For robots to achieve widespread adoption and incorporation into everyday tasks, it is critical that non-experts be capable of imparting knowledge to them through familiar means.

LfD has established itself as a valuable point of interest within the human-robot interaction (HRI) research community, particularly with respect to the intuitive interface it provides users that have never interacted with robots. Members of the HRI community have performed user studies on various LfD techniques, and continue to present results valuable to interaction designers [4]–[7]. Results from applications of these works demonstrate that the effectiveness of each training method is dependent upon factors inherent to the nature of the skill being taught as well as the comfort of the user with manipulating the robot itself.

Some systems have been developed to navigate the spectrum between guidance-based learning (such as LfD) and exploration-based learning (like reinforcement learning, or RL) to operate more efficiently in environments where human collaboration is expected. Systems that leverage socially guided exploration are equipped to learn from non-experts during task execution through familiar forms of social learning. Including a human in the loop for skill acquisition greatly increases the potential value of time spent training. One technique that particularly takes advantage of humans in the loop is the combination of socially guided machine learning and active learning [8], [9].

Active learning is a method designed to accelerate time to skill acquisition. By enabling the learner to ask questions of its teachers, the learner can guide instructors to fill the most critical gaps in its knowledge with training data. Work by Cakmak and Lopes details an algorithm capable of identifying the most desirable areas of training data through analysis within sequential decision problems [10]. This result is applied in a user study, showing gains in performance by agents trained by humans who leveraged the agent's capacity for active learning over those that did not. Designing active learning into systems that interact with non-experts requires special care, as designing for certain query types has been shown to influence user perception [11]. Active learning remains an open topic of interest within HRI. Guidelines for indicating appropriate types of agent queries within various types of scenarios are only beginning to emerge, and will become increasingly important as the complexity of collaborative systems increases.

For robots working collaboratively and sharing goals with humans, having relatable mental representations of skills and tasks contributes to behaviors that are more comprehensible, facilitating the completion of complex tasks. Internal task decomposition is at the core of this mental model. Various existing algorithms are capable of decomposing tasks into representations agreeable for the internal processing of an autonomous agent [12]–[15]. Collaborating with humans introduces a preference for representations that are inherently comprehensible to people, avoiding the introduction of any unnecessary cognitive load when properly implemented [16], [17]. Recent work in this area has begun yielding effective results, including an action segmentation algorithm inspired by human-like methods of segmentation [18]. By utilizing statistical regularities in action sequences performed by humans, a robot observer can discern meaningful options from identifying boundaries between probable groupings of actions. Maintaining and converging upon common mental representations of tasks facilitates interaction design across all levels of abstraction.

A variety of approaches have been studied involving interfaces by which non-programmers can impart plans for fulfilling complex tasks to robots. In addition to the training of primitive skills, enabling non-experts to develop complex interaction behavior for robots is necessary for widespread adoption. In addition to making flexible skill acquisition accessible, compiling higher-level sequences of actions must also be approachable by non-technical users. Rybski et al. introduce an interface by which a human can train a robot verbally through the correlation of location data, agents present, and spoken commands to known skills [19]. The work primarily leverages spoken language to allow non-experts to impart skill sequences to a mobile robot. Breazeal et al. have performed studies utilizing an interface combining verbal communication with joint intention theory [20]. Their system approaches this problem using dynamically interleaved sub-plans to produce coherent teamwork between involved parties, focusing on task representation in terms of goals rather than solely motion trajectories. In lieu of verbal communication, which can quickly become infeasible as task complexity scales, others have developed software interfaces targeted at non-programmers. Glas et al. evaluate this method by way of a user study in which interaction designers were able to use a graphical interface to achieve greater success in programming a complex interaction than groups that did not have access to the interface [21].

III. Hard Problems in Shared-Environment Collaboration

While skill acquisition and engaging non-technical users are clearly essential and relevant fields of research, they alone are not sufficient. The ability to collaborate is reliant upon various physical, perceptual, and social capabilities. For example, a successful teammate often requires a model of its collaborators and the ability to communicate with them. A competent worker requires the ability to execute skills while adapting to a dynamic environment changing outside of its control. Basic operational safety demands a robot with access to a variety of sensors and real-time interpretations of the resulting data inflow, along with knowledge of appropriate actions to respond in kind. Additional software-level precursors such as emotional state recognition, object permanence, and self-evaluation of performance also dictate the potential quality of any possible team activities.

Even if a robot was constructed with dexterous capabilities exceeding that of a typical adult human, a sensor suite capable of producing features useful for complex object recognition, and hardware powerful enough to do real-time complex modeling, shared-environment collaboration in the real world presents a challenging set of research questions. Because of the great technical difficulty and hardware requirements involved, the vast majority of innovations surrounding these questions have only begun to be accomplished by the collaborative robotics (co-robots) and HRI communities.

In this paper, we discuss four of the deepest and most important research questions facing those that seek to build complete systems for human-robot collaboration. The challenges we feel to be most important within this domain are:

- How can co-robots enable and facilitate bi-directional intent recognition?
- How does a co-robot know what roles to take when working with human teammates?

- How does a co-robot know when to trade off task execution optimality for co-worker preferences?
- How does a co-robot self-evaluate during live collaboration?

We seek to motivate investigation into these questions with representative examples from the domains of construction and cooking. Though many human-robot teams will comprise a many-to-many relationship, we focus on one-to-one teams as a scalable baseline for highlighting fundamental sub-questions that need be addressed within each of these four major areas.

From the domain of construction, we consider the task of assembling a structure such as a small animal shelter. The assembly of this structure requires two plastic storage containers, polystyrene foam insulation, and straw. It is constructed by cutting the insulation to shape, placing the insulation inside one container, nesting the second container within the first, and placing straw in the inner container. The construction is completed by placing and sealing the lids on the containers and cutting an access hole in one side. This task embodies many challenges present in collaborative task execution, including division of responsibilities, fine motor control, object manipulation, and potentially dangerous actions (the cutting of materials).

From the domain of cooking tasks, we consider making cherry gelato. To make a gelato base, a chef requires milk, heavy cream, egg yolks, and sugar. The difficulties inherent to gelato production include the need for precise control of temperatures and consistencies. Milk and cream must be mixed and heated, after which egg yolks and sugar are mixed in and heated to a particular consistency (e.g., just coating a spoon). Finally, the mixture is strained into a bowl and chilled until ready to freeze. For the added flavoring, cherries must be pitted and chopped, mixed with sugar and lemon juice, heated, and pureed before being combined with the base mixture. This task entails precise measurements and observations, manipulation of materials with delicate or complex physical properties, and rigid temporal constraints. Given the complexity of the subtasks, many cooking tasks inherently have steps best suited to humans and steps best suited to robots [22].

A. How can co-robots enable and facilitate bi-directional intent recognition?

Humans tend to communicate intent through a variety of methods during collaborations, enabling co-workers to predict each others' future actions and plan around them. This communication can occur both verbally and non-verbally, potentially spanning multiple levels of abstraction. Teams that communicate implicitly improve performance under stresses caused by temporal constraints or uncertainty by acting in anticipation of teammates' actions [23]. For example, suppose two humans (X and Y) are constructing the animal shelter from the construction example above. If X requires use of the cutting tool, X's gaze may orient towards the desired object. Y may be capable of anticipating X's need based on prior observed action history, or may require more information to realize that assistance is being requested. Y can actively calibrate his intention recognition by holding the perceived object of desire in the air while focusing attention on X, as X is also likely to provide a response confirming or disconfirming Y's interpretation. This interaction communicates the same information across a multiple channels, including joint attention, object manipulation, and socially meaningful gestures.

Active engagement within intention recognition is a powerful tool for a robot to leverage, with the potential to greatly enhance the speed and quality of learned collaboration [24]. Collaborative robots must have robust classifiers capable of determining human intent, especially in situations where two agents' skills require coordination, like object transfer. By conveying intent at a high level, such as that of role selection, a co-robot can guide human teammates to choose roles with non-conflicting subtasks for greater overall efficiency. At a finer level, such as conveying intention within individual skills, the ability to broadcast one's goals or intended motion paths will lead to fewer instances of turn-taking behavior and conflicts over occupancy of shared spaces. Intention recognition is a dynamically calibrating process, where each agent plays an active role of synchronizing communication channels and establishing expectations.

Identifying intent is especially important for potentially dangerous tasks, such as cutting an opening into the animal shelter. Tasks that endanger either humans or robots necessitate the prior conveyance of intention for basic operational safety, beyond any safeguards inherent to the motor control sequence itself. From an HRI standpoint, collaborating humans' perceptions of danger are as important to consider as any actual danger during robot operation.

Beyond safety concerns, anticipation of where objects will be placed or moved while in use directly affects the possible actions or available space usage of collaborators. As available space is a precious resource within shared-space collaborative task execution, the importance of communicating one's intent grows inversely proportionally to the amount of available working space usable by the team. In the context of the cooking example, a situation where the active working space is presumably a small area such as a kitchen, communicating one's intended motion path is critically important. Fragile or loose materials that may become dangerous under improper handling (e.g., hot liquids) must be transported with care. This requirement can be satisfied without sacrificing advantages afforded by the parallel nature of the task, but only if team communication is effective enough to enable it.

Early research in bi-directional intention recognition has focused on pre-execution communication and anticipatory motion. Existing approaches rely upon attempting to characterize significant features used in human-to-human collaborative activities, such as handover tasks [25], [26]. These lab-based studies measured the importance of various features on detecting an intent to handoff an object, including partner orientation, gaze direction, hand occupancy, social gestures, and partner distance. These results were used to build a classifier with human-interpretable rules capable of recognizing handover intention within their coded dataset. It is important to follow up on these studies with deployed collaborative systems to understand how bi-directional intention recognition can function in more stochastic environments.

Some questions that remain unanswered with regard to intention conveyance include: How can effective non-verbal cues be generated for learned skills? How can a robot leverage channels of communication that humans understand, despite dissimilar physical forms or capabilities? How can a robot ensure the comfort and safety of its teammates throughout its operation without compromising the ability to occupy the same working environment? How can a robot identify human intent for skill execution, and how can this information be used by the robot to better select actions to execute?

B. How does a co-robot know what roles to take when working with human teammates?

When participating as a member of a team, fluency of operation can be achieved by determining a division of labor within a task and assigning the various divisions to roles. Each team member's role dictates their expected actions, simplifying the task of predicting their intent or future movement. Furthermore, a division of labor may be more efficient than single-stream task execution, especially if roles are chosen with collaborators' skills in mind. Humans are capable of decomposing tasks into multi-role, multi-collaborator endeavors, and synchronizing these decompositions between them with minimal verbal communication. This decomposition has a strong possibility of being ambiguous, requiring a wealth of contextual knowledge to disambiguate, which is unlikely to be explicitly available to a robot. Being able to generate potential role divisions from a task, synchronizing them with co-workers, and properly selecting a role based on the demands (both social and practical) of the team greatly boosts a co-robot's ability to function as an effective member of a team.

Role assignment is not a static, one-time activity. Pre-activity role determination is important for setting initial roles and expectations, but real collaboration can involve role trading and overlap, the likelihood of this increasing with task complexity. Co-workers often bridge the boundaries between roles to assist each other, even when not explicitly specified in the duties associated with an assigned role. Recognizing implicit communication and anticipating the needs or actions of co-workers contributes to fluent collaboration, increasing both the objective and perceived value of the robot [27]. Roles may change or adjust throughout a task's execution, requiring the capacity to be sensitive to others' preferences, the ability to evaluate one's own and others' skill proficiencies, the ability to evaluate temporal constraints of subtasks between and within roles, and the potential to build action policies within given safety constraints. These components must all become part of the role selection mechanism, evolving over time as more about teammate preferences are revealed.

For a team to achieve fluency of collaboration, learning patterns and appropriate reactions from task repetition is essential. Adaptive planning systems that are capable of modeling and adjusting to human behaviors with social understanding are only beginning to be developed and tested in the real world. Important work yet to be addressed within this space includes the handling of unfulfilled commitments by teammates and handling execution-time variation in ordering constraints [28], [29]. A complete system capable of adapting to team preferences and behaviors requires many other components. In addition to the abilities listed above, a complete system requires the ability to identify roles chosen by teammates through participation and live observation, as well as the ability to estimate the impact of one's actions on others given knowledge about their goals.

Task decomposition algorithms that operate on input demonstrations have achieved success in learning execution policies from low repetitions of demonstrations of skill sequences [12], [30]. They provide a great benefit through reducing the complexity of potentially intractably large problem spaces. This is accomplished by limiting the number of relevant input dimensions according to the demands of particular segments of a task. However, these task decompositions are typically optimized for internal use and are not meant to be interpreted by the user. In collaborative domains, it is important to be able to communicate one's understanding of a task in a manner comprehensible to one's teammates, while maintaining the ability to actively participate. Reconciling this notion with the output from automated task decompositions presents another aspect of the immense challenge inherent to comprehensibly representing one's mental state to another agent.

The objective of a robot teammate for any collaborative activity should be to reduce the workload, either physically or cognitively, of fellow teammates. A co-robot should be capable of identifying divisions of responsibilities within tasks, synchronizing its plan with its teammates, and self-selecting appropriate roles to assume during execution. These behaviors must be adaptive and flexible to both plan and role-assignment changes during execution.

Live role selection requires answers to the questions: How can a robot learn decompositions of complex tasks? How can a robot synchronize these decompositions with human teammates? How can a robot reduce the dimensionality of required skills such that they become tractable to represent or execute? Given a task decomposition and list of roles associated with actions within said decomposition, how does a robot know which role(s) to assume while minimizing conflict with co-workers? How can a robot transfer previously observed team dynamics to new tasks? How can a robot communicate its internal task decomposition in an understandable manner?

C. How does a co-robot know when to trade-off task execution optimality for co-worker preferences?

When operating in mixed human/robot teams, it is inevitable that tasks will require the use of skills that certain teammates are better at than others. Likewise, there will be situations in which certain subtasks should be handled explicitly by a robot and situations where subtasks are best suited to human execution. For example, dangerous or repetitive tasks—cutting through plastic or pitting 100 cherries—are good candidates for robotic execution, while tasks that require human-human communication or high-level decision making —getting size specifications for a shelter or measuring the consistency of gelato—are ideal for human execution. Evaluating co-worker proficiency and recognizing these situations presents a difficult research challenge [31]. Even if a system were capable of this, however, the question of when and how this information should be leveraged remains.

One approach is to develop a plan over several iterations in concert with the robot teammate(s) to be used. This allows the plan to evolve naturally with all collaborators present in the process, engendering trust between workers [32]. In the future, collaborative robotics will need to enable robots to join existing teams without incurring the high costs associated with disrupting existing dynamics and redefining roles. This question introduces many others, perhaps most notably: How does one balance between preferences of co-workers, the necessities of the assigned task, and the performance criteria with which the group is being judged?

Further adding to the difficulty of this question, the performance criteria may be unknown, varying by team and task. An agent must be able to successfully balance time, monetary cost, energy expended, and even emotional costs incurred by someone doing a task against their preferences. Psychological considerations may also be taken into account, including methods of handling co-worker disengagement or dissatisfaction. This evaluation function may also be subject to an external authority dictating constraints, which may occasionally be in conflict with teammate preferences. A complete collaborative system should be able to determine which trade-offs are most important and what is the best possible decision considering the team as a whole. This not only requires robust models of each teammate's intent as covered previously, but also the ability to take measurements of teammate performance, each of which evolves over time.

As a participant in the animal shelter construction example, a robot may be more precise when cutting insulation than humans, but whether that role is a desirable choice may depend on the preferences of the human. Within the cooking example, it may be optimal for the robot to perform steps that have strict monitoring requirements such as heating milk or custard, but a human may not trust the robot to perform the task safely or with proficiency. Some tasks may be too complex or costly to perform full trials of for demonstrative purposes. How can a robot instill trust in co-workers without the luxury of demonstration as proof?

Pertinent questions within the concept of optimality versus preferences tradeoff include: How can a robot adapt its role selection to optimally mesh with an already-established human team? When should a robot collaborator object to an observed role distribution? How can a robot learn which subtasks it should handle rather than humans? How does one quantify ongoing co-worker satisfaction? How can this data be used to impact a robot's role selection decisions? How can a robot mediate disputes when co-worker preferences are at odds with optimal roles?

D. How does a co-robot self-evaluate during live collaboration?

Interacting with multiple agents in a shared task presents opportunities to evaluate personal proficiencies. While individual skills often include measures of success that can be evaluated during execution, assessing one's proficiency at being an effective and useful team member can be less straightforward. Determining this quality requires a shared mental model of the task, knowledge of co-workers' roles and responsibilities, and accurate estimates of expected overall task progress at given times. Even still, determining the appropriate metrics and acquiring reasonable real-time measurements presents a substantial research challenge. Once acquired, this information can be used to reinforce role selections, refine skill execution choices, and evaluate novel task decompositions.

A collaborative robot must use some form of contingency detection to measure the impact it has on its teammates at small timescales. Contingency detection is the detection of a change in an agent's behavior within specified time bounds of another agent's action. This phenomenon has been studied within HRI for several reasons, including its utility in understanding differences between oneself and others [33] and in general for classifying responses (or non-responses) to actions [34]. Some planners have begun to use these behavioral or temporal fluctuations to increase the fluidity of team operation, adapting to perceived co-worker preferences in task execution [35].

Individuals are capable of using a wide variety of metrics to gauge success. Beyond mere evaluation of whether the group completed the assigned task, self-evaluation can include questions measuring personal performance or growth as well. A co-robot must be able to judge whether it performed its assigned tasks properly; for instance, it must be able to evaluate not just whether it has made gelato, but whether the gelato actually tastes good. Success may also partially be defined as self-improvement with regard to various subskills used throughout the task; for example, a robot may have developed its insulation cutting skill to be more efficient during construction. Building cohesive teams requires pairing workers who value each others' contributions, so co-worker satisfaction and satisfaction levels of authority figures should also comprise a component of a comprehensive self-evaluation function.

During task execution, humans are capable of evaluating a teammate's progress executing a skill against their own perceived ability to perform the same skill. A co-robot with this capability introduces new potential issues related to team management. How does a co-robot know when to switch off of its current task, either due to a recognized inability to perform it efficiently or because it has discovered it can perform a different required task with better results? At what performance threshold does it become appropriate to intercede with an inefficient human worker for the greater good of the team? What are the psychological implications of a robot co-worker suggesting a teammate is not performing at his or her potential?

Other pertinent questions within collaborator self-evaluation include: How does a robot detect if it is performing in line with its teammates' expectations? How can a robot evaluate its proficiency at performing its assigned role within the context of the team at large? How can contingency detection scale to scenarios involving three or more agents while disambiguating false positives for teammate actions? What are appropriate actions for a robot to take if it detects it is underperforming to best maintain team cohesion? When should a robot ask for help with a task if it means interrupting an otherwise occupied teammate?

IV. Conclusion

Human-robot collaboration promises to revolutionize the way people work, following from the world-altering changes of industrial robotics. In this paper, we have presented four questions that highlight topic areas within collaborative robotics. Our discussion of each question contains a multitude of research ideas and critical points to be addressed by future works studying human-robot teaming. Within each topic area, we have provided examples of current work and as well as examples of how specific collaborative tasks could benefit from answers to each of the four questions. Our aim is to call attention to what we perceive to be some of the most critically underserved and promising research opportunities within the rich field of human-robot collaboration. We seek to inspire researchers to look at questions essential for building complete systems capable of fluid, flexible, and natural collaboration between arbitrary collections of skilled agents.

ACKNOWLEDGMENT

This material is based upon work supported by Office of Naval Research grant #N00014-12-1-0822 ("Social Hierarchical Learning") and grants from the National Science Foundation under contracts No. 1139078, No. 1117801, and No. 0835767.

REFERENCES

[1] R. Sutton, D. Precup, S. Singh *et al.*, "Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning," *Artificial intelligence*, vol. 112, no. 1, pp. 181–211, 1999.

[2] V. Groom and C. Nass, "Can robots be teammates? benchmarks in humanrobot teams," *Interaction Studies*, vol. 8, no. 3, pp. 483–500, 2007.

[3] B. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009.

[4] J. Sweeney and R. Grupen, "A model of shared grasp affordances from demonstration," in *Humanoid Robots, 2007 7th IEEE-RAS International Conference on*. IEEE, 2007, pp. 27–35.

[5] S. Calinon and A. Billard, "Incremental learning of gestures by imitation in a humanoid robot," in *Proceedings of the ACM/IEEE international conference on Human-robot interaction*. ACM, 2007, pp. 255–262.

[6] C. Atkeson and S. Schaal, "Robot learning from demonstration," in *International Conference on Machine Learning*, 1997, pp. 11–73.

[7] B. Akgun, M. Cakmak, J. Yoo, and A. Thomaz, "Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective," in *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 2012, pp. 391–398.

[8] A. Thomaz and C. Breazeal, "Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance," in *Proceedings of the National Conference on Artificial Intelligence*, vol. 21, no. 1. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2006, p. 1000.

[9] C. Breazeal and A. Thomaz, "Learning from human teachers with socially guided exploration," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*. IEEE, 2008, pp. 3539–3544.

[10] M. Cakmak and M. Lopes, "Algorithmic and human teaching of sequential decision tasks," in *National Conference on Artificial Intelligence (AAAI12), Toronto, Canada*, 2012.

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

[12] G. Konidaris, S. Kuindersma, A. Barto, and R. Grupen, "Constructing skill trees for reinforcement learning agents from demonstration trajectories," *Advances in neural information processing systems*, vol. 23, pp. 1162–1170, 2010.

[13] P. Stone and M. Veloso, "Task decomposition and dynamic role assignment for real-time strategic teamwork," *Intelligent Agents V: Agents Theories, Architectures, and Languages*, pp. 293–308, 1999.

[14] G. Konidaris and A. Barto, "Skill discovery in continuous reinforcement learning domains using skill chaining," *Advances in Neural Information Processing Systems*, vol. 22, pp. 1015–1023, 2009.

[15] S. Calinon, F. Guenter, and A. Billard, "On learning, representing, and generalizing a task in a humanoid robot," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 37, no. 2, pp. 286–298, 2007.

[16] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "Situation awareness, mental workload, and trust in automation: Viable, empirically supported cognitive engineering constructs," *Journal of Cognitive Engineering and Decision Making*, vol. 2, no. 2, pp. 140–160, 2008.

[17] P. S. Tsang and M. A. Vidulich, "Mental workload and situation awareness," *Handbook of Human Factors and Ergonomics, Third Edition*, pp. 243–268, 2006.

[18] J. Shim and A. Thomaz, "Human-like action segmentation for option learning," in *RO-MAN, 2011 IEEE*. IEEE, 2011, pp. 455–460.

[19] P. Rybski, K. Yoon, J. Stolarz, and M. Veloso, "Interactive robot task training through dialog and demonstration," in *Human-Robot Interaction (HRI), 2007 2nd ACM/IEEE International Conference on*. IEEE, 2007, pp. 49–56.

[20] C. Breazeal, G. Hoffman, and A. Lockerd, "Teaching and working with robots as a collaboration," in *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 3*. IEEE Computer Society, 2004, pp. 1030–1037.

[21] D. Glas, S. Satake, T. Kanda, and N. Hagita, "An interaction design framework for social robots," *Robotics: Science and Systems VII*, p. 89, 2012.

[22] M. Bollini, J. Barry, and D. Rus, "Bakebot: Baking cookies with the pr2," in *The PR2 Workshop: Results, Challenges and Lessons Learned in Advancing Robots with a Common Platform, IROS*, 2011.

[23] J. Shah and C. Breazeal, "An empirical analysis of team coordination behaviors and action planning with application to human–robot teaming," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 52, no. 2, pp. 234–245, 2010.

[24] C. Crick and B. Scassellati, "Intention-based robot control in social games," in *Proceedings of the Cognitive Society Annual Meeting*, 2009.

[25] K. Strabala, M. Lee, A. Dragan, J. Forlizzi, and S. Srinivasa, "Learning the communication of intent prior to physical collaboration," *Robot and human interactive communication*, pp. 968–973, 2012.

[26] K. Strabala, M. Lee, A. Dragan, J. Forlizzi, S. Srinivasa, M. Cakmak, and V. Micelli, "Towards seamless human-robot handovers," *Journal of Human-Robot Interaction*, vol. 1, no. 1, pp. 1–23, 2013.

[27] G. Hoffman and C. Breazeal, "Effects of anticipatory action on human-robot teamwork: Efficiency, fluency, and perception of team," in *Human-Robot Interaction (HRI), 2007 2nd ACM/IEEE International Conference on*. IEEE, 2007, pp. 1–8.

[28] 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.

[29] J. Shah, "Fluid coordination of human-robot teams," Ph.D. dissertation, Massachusetts Institute of Technology, 2011.

[30] L. Cobo, C. Isbell Jr, and A. Thomaz, "Automatic task decomposition and state abstraction from demonstration," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2012, pp. 483–490.

[31] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 30, no. 3, pp. 286–297, 2000.

[32] S. Nikolaidis and J. Shah, "Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy," 2013.

[33] K. Gold and B. Scassellati, "Learning acceptable windows of contingency," *Connection Science*, vol. 18, no. 2, pp. 217–228, 2006.

[34] J. Lee, J. Kiser, A. Bobick, and A. Thomaz, "Vision-based contingency detection," in *Proceedings of the 6th international conference on Human-robot interaction*. ACM, 2011, pp. 297–304.

[35] R. Wilcox, S. Nikolaidis, and J. Shah, "Optimization of temporal dynamics for adaptive human-robot interaction in assembly manufacturing," in *Proc. RSS*, 2012.