Human-Robot Teaming using Shared Mental Models

Stefanos Nikolaidis*
Massachusetts Institute of Technology
snikol@mit.edu

Julie Shah†
Massachusetts Institute of Technology
julie_a_shah@csail.mit.edu

ABSTRACT

Robots are increasingly introduced to work in concert with people in high-intensity domains, such as manufacturing, space exploration and hazardous environments. Tasks in these domains are often well-defined, but involve complex coordination under constraints and are performed under time pressure. Although numerous studies on human teamwork and coordination in these settings, very little prior work exists on applying these models to human-robot interaction. In this paper we propose a methodology for applying prior art in Shared Mental Models (SMMs) to promote effective human-robot teaming. SMMs are measurable models developed among team members prior to task execution and are strongly correlated to team performance.

Keywords

human robot collaboration, human robot interaction, teamwork, mental models, interactive robots, team performance

1. INTRODUCTION

We propose a novel framework that uses insight from prior art in human team coordination and shared mental models to increase the performance of human-robot teams collaboratively executing complex tasks. Shared Mental Models (SMMs) [4] are measurable models developed among team members prior to task execution and are strongly correlated to team performance. Although numerous studies have modeled the performance-linked characteristics of SMMs in human team coordination, very little prior work exists on applying these models to a human-robot interaction framework. We propose that valuable insights can be drawn from these works. For instance, a study evaluating teamwork in flight crews [15] has shown that teams with accurate but different shared mental models among team members perform worse than teams having less accurate but common models. Applying this insight to human-robot teaming leads to a hypothesis that, to promote effective teamwork, a robot must execute a task plan that is similar to the human partner's mental model of the execution. In this paper, we outline a framework that leverages methods from human factors engineering to promote the development of teaming models that are shared across human and robot team members. The human factors community has developed and validated a widely used set of techniques for eliciting, quantitatively evaluating, and strengthening SMMs for human teamwork [13, 20, 14]. Studies of military tactical operations and aviation crews show that improved team performance and reduction of errors are related to both the quality and similarity of the team members' shared mental models [15, 16]. High quality shared mental models facilitate the use of implicit communications, which are documented to play a primary role in effective human team coordination [22, 24]. These findings have potentially important implications for designing safe and effective human interaction with robots in safety-critical domains.

The proposed approach leverages this body of work and our understanding of cognitive behavioral psychology to develop new computational methods that facilitate (1) interactive planning between human and robot co-workers, (2) assessment of team models, and (3) promotion of effective team interaction. The expectation is that this approach will yield quantitative improvements in human-robot team safety and performance, as demonstrated for human teamwork in the military, aviation, and medical domains.

Our approach distinguishes two phases of human-robot teamwork that mirror the standard techniques for promoting effective human team coordination [24, 27, 13] a **planning phase**, and a **task execution phase**. The purpose of the planning phase is to derive a high-quality teamwork model that provides a perceptual common ground to predict coworkers' actions, and integrate predicted effects of the humans' and robots' actions. The proposed methodology investigates a two-part planning phase: first, the robot independently derives a team model for the task by observing a team of expert human workers. The robot then executes an interactive planning and cross-training process with a human co-worker, to iteratively refine and converge the team model. The key hypothesis to be tested is that human-robot performance at the execution phase can be improved by achieving similar and accurate Shared Mental Models for the human and robot in the planning phase.

*Graduate Student

†Boeing Assistant Professor, Department of Aeronautics and Astronautics

2. APPROACH AND METHODOLOGY

We outline our proposed methodology in the following four steps:

(1) The robot derives an initial teaming model from observation of coordinated manual work performed by two or more expert human workers. The teaming model is described as a Partially Observable Markov Decision Process (POMDP). This step is discussed further in Section 2.1.

(2) Next, the robot and a new worker perform structured interactive planning in a computer simulation environment. We focus on computationally emulating cross-training, which is defined as "an instructional strategy in which each team member is trained in the duties of his or her teammates" [27] (Section 2.2.1). Through this cross-training process, the robot iteratively refines its teaming model until it converges with the human workers' mental models.

(3) We assess the quality and convergence of the learnt model.

(4) Finally, we apply the converged teaming model to perform robot control on the production line or operational environment. We will evaluate the validity of our framework by (a) empirically investigating the similarity of the robot's teaming model and the elicited human mental models using standard techniques [5], and (b) quantitatively assessing human-robot team performance at the execution phase. Each step is described further in the following subsections.

2.1 Deriving a team mental model

The literature presents various definitions for *team mental model* [11]. In our work we use the definition in Marks et al. [13], that mental models contain "the content and organization of interrole knowledge held by team members within a performance setting ... and ... contain procedural knowledge about how team members should work together on a task within a given task domain, including information about who should do what at particular points in time". We address the challenge of deriving a computational representation of a teaming model as a statistical inference problem. Whereas some task requirements and specifications are provided explicitly a priori, the robot develops an accurate computational model of its actions with respect to the action of its human co-worker, by learning from observation of two expert human workers working together. The implicit nature of mental models, the uncertainty in the observations of actions and the prevalence of information-seeking behavior in human team interactions makes a Partially Observable Markov Decision Processes (POMDP) representation attractive [10].

2.1.1 POMDP Formulation

We explore a POMDP representation [10] as the computational framework for encoding the teaming model. A partially observable Markov decision process can be described as a tuple $\{S, A, T, R, \Omega, O\}$, where

- $\{S$, A, T, and $R\}$ describe a Markov decision process;
- Ω is a finite set of observations the agent can experience of its world; and
- $O : S \times A \longrightarrow \Pi(\Omega)$ is the *observation function*, which gives, for each action and resulting state, a probability distribution over possible observations (we write $O(s', a, o)$ for the probability of making observation o given that the agent took action a and landed in state s').

We propose that mental model features can be expressed as a sequence of actions (policy π) over states. The robot uses the POMDP model as the framework for learning an action strategy (policy) based on a sequence of recorded demonstrations of two human experts executing a task. In order to capture a wide range of mental model features, the POMDP model's set of actions includes task-based actions (pick-and-place) and interaction-based actions (gestures, motions, posture, and simple voice commands). As state we use a vector of variables. Augmenting the state in a POMDP framework has been used in previous literature for the purpose of assisting patients with dementia [7] and elderly with mild cognitive and physical impairments [17].The state vector S includes:

- *planning system variables* for parametrized representations of the relative pose between co-workers,
- *activity status variables* to encode actionable sub-tasks (e.g. "assemble part A," "connect part A to part B") and partial ordering constraints among sub-tasks,
- *interactive variables* including gestural and verbal communicative acts and physical interactions such as joint manipulation of objects, and
- *mental model variables*, which include parameterized encodings of worker responsiveness, interdependence, and role (e.g. leader/initiator, follower/responder).

Observations Ω are derived from human motion and speech recognizers, which classify task-based motions and speech utterances in real-time. We use a Phasespace optical motion tracking system to provide real-time ground-truth on the human worker's position, posture and movements. We capture key features of shared mental models through elicitation and adaptation of the structure of the POMDP, the transition probabilities T relating state to action. The probabilities are initialized based on a priori knowledge of the task instructions and constraints. Structure is then adapted based on observation of several task execution sequences from two human experts, and through iterative reinforcement learning during the interactive planning phase. We use the task specifications to set rewards R within the POMDP model to a large positive value for each successful completion of a subtask. The robot's optimal policy may be computed using the PoliCy-Contingent Abstraction (PolCA) algorithm [19], in which each task is mapped into a hierarchy of subtasks, and solved separately. This algorithm performs well for hierarchically structured problems, as solving for N subtasks is usually computationally more efficient than solving task which is N times large.

2.2 Robot and human information sharing

It has been shown in previous research that in cases where effective communication is a key component for the mission success, team members with similar mental models perform better than teams with a more accurate but less similar mental models [16]. Therefore, even if the mental model learnt by observation of the two human experts is accurate, the robot needs to adapt this model when asked to work with a new human partner. The goal then becomes for the human-robot team to develop a shared mental model. The mental model adaptation process needed is often described as “cross-training” in literature [13]. We plan to emulate the cross-training process among human team-members by having human and robot train together at a virtual environment, until they reach a shared mental model. In the following subsection we describe the human-robot cross-training paradigm, which is similar to the cross-training process for human teams.

2.2.1 Cross-training

There are three types of cross-training [3]: a) Positional clarification (b) positional modeling and c) position rotation. Findings suggest that position rotation, which is defined as “learning interpositional information by switching work roles”, is the most strongly correlated to improvement in team performance [13]. We plan to implement positional rotation in the human-robot team scenario by having the human worker switch roles with the robot at a virtual environment. The robot then will refine the computed POMDP model according to the new samples. We view this as an iterative reinforcement learning process, where human and robot adjust the interdependent action sequences to maximize a reward. We expect the cross-training to result in human and robot converging to a shared mental model, subsequently on improved human-robot team performance.

2.3 Evaluation

2.3.1 Mental Model Elicitation

After cross-training, we assess the similarity of mental models of the human and robot. Our approach generalizes from standard techniques for eliciting and evaluating shared teaming models for knowledge-based tasks. Following prior work that quantitatively assesses mental model similarity for human teams [14], we use a variant of the Pathfinder program [5] to derive relational and similarity measures for SMMs in human-robot teams.

2.3.2 Animation-Based Pathfinder for Human-Robot Teaming

We describe prior work [28] assessing team-interaction models for military tactical teams as an illustrative example for how Shared Mental Model elicitation techniques are applied. First, subject matter experts conduct a comprehensive team task analysis and identify task-related concepts across the team member roles. Team members are asked to provide ratings of the relations among critical task concepts. Some task-related concepts are closely interrelated, while others are weakly coupled. For example, identifying the enemy, performed by the radar specialist, and selecting the appropriate weapon for that enemy, performed by the gunner, are quite interrelated. Adjusting altitude is highly related to escaping enemy attacks, though weakly related to selecting waypoints. Team and task-related knowledge is documented using a relational matrix and a Likert-type scale that ranges from 1 (not related) to 9 (very related). These data are input to the Pathfinder program [8], a computerized networking technique that derives network structure based on the perceived relatedness among task-related concepts. Pairwise relations between concepts are represented by links and spatial proximity. Pathfinder produces a similarity index, the C index, reflecting the similarity among each pair of team members' relational matrices.

Pathfinder-like systems, including variants that encode time-evolved behaviors, have been applied widely and have been empirically validated to correlate to team performance outcomes [12]. However, a Pathfinder-like network-analysis approach to eliciting and assessing team models has not yet been developed for physical interaction tasks. Also, networks produced by Pathfinder-like systems are not appropriate to model and drive robot control. The proposed approach will apply the spirit of the Pathfinder system to new computational techniques for representing teaming models for physical interaction tasks.

Interactive planning and model convergence will be performed using a virtual, animated simulation environment and a structured view-show-modify process to prototype different human-robot interactions. We plan to add Pathfinder-like similarity assessments in the cross-training process, integrating them as embedded tasks in the simulation environment. Relational information about sequencing and timing of actions, communication, and strategies for resource utilization will be portrayed and elicited, for example, through animations of human-robot interactions.

2.3.3 Human-Robot Task Execution

When the planning phase is finished, human and robot will execute the required task in real-time. As the human and robot will share a common model, we anticipate better team coordination and performance for the task. Furthermore, we anticipate the robot will execute anticipatory actions assisting the human using the refinements of the computational model made at the cross-training phase (Section 2.2.1). We base this hypothesis on recent findings that robot anticipatory actions in a human-robot team setting lead to significant improvement in task efficiency and in the perceived commitment of the robot to the team [8]. We will evaluate our framework by comparing human-robot team performance after cross-training, with the performance of the same team without the cross-training phase. For the performance evaluation we plan on using the three fluency metrics proposed in [8]: concurrent motion, human idle time, and time between human and robot action. Furthermore, we will perform a user study to show the effect of the cross-training process on the human worker experience.

3. APPLICATION STUDY: ASSEMBLY MANUFACTURING

Traditionally, industrial robots in manufacturing and assembly perform work in isolation from people. When this is not possible, the work is done manually. In the automotive industry the result is a split factory; on one side all robots, on the other side all people. The impact of this imposed dichotomy is more pronounced in the aerospace manufacturing industry. Assembly of large commercial airplanes is primarily manual work. Tight integration and variability in the build process makes it difficult to physically isolate elements for robotic-only work without significant detriment to efficiency and workflow.

We envision a new class of manufacturing processes that achieve significant economic and ergonomic benefit through robotic assistance in manual processes. For example, mechanics in aircraft assembly spend a significant portion of their time retrieving and staging tools and parts for each job. A robotic assistant can provide productivity benefit by performing these non-value-added tasks for the worker.

Programming a robot to work in a team with a person is not currently possible using standard industrial robot teach pendants. Mechanisms exist to program multi-robot cells, but these require repeatable, coordinated movements with preprogrammed synchronization points. These techniques are not applicable to programming human-robot teamwork, as workflow changes from worker to worker based on personal preferences and technique. Substantial difficulties arise in attempting to capture convergent team behavior using traditional programming scripts. Next, we present an example assembly task that illustrates the potential for adaptive human-robot collaboration in assembly.

3.1 Motivating Example: Robot Assistant for an Assembly Mechanic

We aim to develop a capability that supports efficient and productive interaction between a worker and a robotic assistant, such as the ABB FRIDA robot shown in Fig. 1. Although important aspects like tolerances and completion times are well defined, many details of assembly tasks are left largely up to the mechanic. Assembly of airplane spars is one example of a manual process where mechanics develop highly individualized styles for performing the task. Fig. 2 shows a mechanic assembling a spar. The spar is composed of two pieces that must be physically manipulated into alignment. After alignment, wet sealed bolts are hammered into pre-drilled holes and are fastened with collars. Excess sealant is removed, and the collars are re-torqued to final specifications. Sequencing of these tasks is flexible, subject to the constraint that the sealant is applied within a specified amount of time after opening it.

Figure 1: ABB FRIDA robot acting as a robotic assistant

A robot such as FRIDA can assist a mechanic by picking bolts and fasteners from a singulator, rotating them in front of a stationary sealant end-effector, and inserting them into the bores. This would allow the mechanic to focus on wiping sealant, hammering the bolts, and placing and torquing the collars. This division of labor would provide productivity benefit through parallelization of tasks.

Figure 2: Spar assembly is a manual process that could be improved by a robotic assistant

Our aim is to enable a robotic assistant to adapt to person-specific workflow patterns. If most mechanics like to hammer all bolts before torquing collars, the robot will support this approach by placing all bolts in a pattern that anticipates the mechanic's actions. When the robot is paired with one of the mechanics that instead prefers to hammer and torque the collar for each bolt as it is placed, the robot will quickly perceive this difference and reoptimize its schedule to converge on a turn-taking pattern with the mechanic. The robot will adapt according to the mechanic's preferences, subject to the constraint that the sealant is utilized within the specified window.

3.2 Related Work in Collaborative Industrial Robots

Recent research in human-robot collaboration for manufacturing has primarily focused on the safety aspects of working with industrial robots. The National Institute of Standards and Technology (NIST) has recently begun a project to develop a technical specification for the safe operation of collaborative industrial robots [25]. This includes the use of advanced sensors to detect human positions and adjust robot speed based on the separation distance between the human and the robot. There is also a significant research effort in Germany, the JAHIR CoTeSys project, aimed at enabling safe physical human-robot co-work [12, 6, 28]. This research effort is safety-focused and is aimed at performing real-time sensing and tracking of human co-workers, performing dynamic collision avoidance and adapting the robot's operational position and posture. Related efforts in Germany have investigated gesture and interaction-based mechanisms for flexible teaching of industrial robots (NEUROS and MORPHA projects) [23, 2]. This work leverages recent advances in the field of learning from demonstration [9, 18, 26, 1], which involves the transfer of assembly skill from a person to a robot.

We are interested in a complementary technical approach to collaborative industrial robots. We are exploring the use of shared mental models in designing high-level, person-specific planning and execution mechanisms that promote predictable, convergent team behavior. This capability is tied to understanding of team members' intent, situational awareness, and risk-sensitive real-time analysis of human behavior. One of the goals of the proposed work is to develop the theoretical foundations and new computational approaches for generalizing human factors techniques to human-robot teaming of physical interaction tasks: we envision a situation where a mechanic trains in a virtual simulation environment to work with a robot the same way a flight crew trains in a simulator to practice effective crew resource management (teamwork) techniques [21].

4. CONCLUSION

In this paper, we outline a framework that leverages methods from human factors engineering to promote the development of teaming models that are shared across human and robot team members. We describe results of prior art in human teaming that lend support to the hypothesis that human-robot performance at the execution phase can be improved by achieving similar and accurate Shared Mental Models for the human and robot in the planning phase. We describe a structured interactive planning process aimed at achieving convergence in teaming models, and present a manual assembly manufacturing task where the proposed methodology will be applied and evaluated.

5. ACKNOWLEDGEMENTS

This work is supported in part by ABB, and is being conducted in collaboration with Thomas Fuhlbrigge, Gregory Rossano, and Carlos Martinez of ABB Inc., USCRC - Mechatronics. We would also like to acknowledge the Onassis Foundation as a sponsor, and thank Mark Gabriel of the Boeing Company for sharing his insight into assembly manufacturing applications, including the spar assembly example.

6. REFERENCES

[1] R. Amit and M. Mataric. Learning movement sequences from demonstration. In *Proc. ICDL*, pages 203–208, 2002.

[2] R. Bischoff, A. Kazi, M. Seyfarth, and K. R. Gmbh. The morpha style guide for icon-based programming. In *In Robot and Human Interactive Communication, 2002. Proceedings. 11th IEEE International Workshop on*, pages 482–487, 2002.

[3] C.-B. J. Blickensderfer E. and S. E. Cross-training and team performance. *Making decisions under stress: Implications for individual and team trainin, American Psychological Association*, pages 299–311, 1998.

[4] N. J. Cooke, E. Salas, and J. A. Cannon-bowers. Human factors, 42, 151-173. measuring team knowledge.

[5] D. Dearholt and R. Schvaneveldt. *Properties of pathfinder networks. In R. Schvaneveldt (Ed .), Pathfinder associative networks: Studies in knowledge organization.* Ablex, Norwood , NJ, 1990.

[6] M. Giuliani, C. Lenz, T. Muller, M. Rickert, and A. Knoll. Design principles for safety in human-robot interaction. *International Journal of Social Robotics*, 2(3):253–274, September 2010.

[7] J. Hoey, P. Poupart, A. v. Bertoldi, T. Craig, C. Boutilier, and A. Mihailidis. Automated handwashing assistance for persons with dementia using video and a partially observable markov decision process. *Comput. Vis. Image Underst.*, 114:503–519, May 2010.

[8] G. Hoffman and C. Breazeal. Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*, HRI '07, pages 1–8, New York, NY, USA, 2007. ACM.

[9] G. E. Hovland, P. Sikka, and B. J. McCarragher. *Skill acquisition from human demonstration using a hidden Markov model*, volume 3, pages 2706–2711. Ieee, 1996.

[10] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artif. Intell.*, 101:99–134, May 1998.

[11] J. Langan-Fox, S. Code, and K. Langfield-Smith. Team mental models: Techniques, methods, and analytic approaches. *Human Factors*, 42(2):242–271, 2000.

[12] C. Lenz, A. Sotzek, T. Roder, H. Radrich, A. Knoll, M. Huber, and S. Glasauer. Human workflow analysis using 3d occupancy grid hand tracking in a human-robot collaboration scenarios. In *In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, San Francisco, USA*, 2011.

[13] M. Marks, M. Sabella, C. Burke, and S. Zaccaro. The impact of cross-training on team effectiveness. *J Appl Psychol*, 87(1):3–13, 2002.

[14] M. A. Marks, S. J. Zaccaro, and J. E. Mathieu. Performance implications of leader briefings and team-interaction training for team adaptation to novel environments. *Journal of Applied Psychology*, 85(6):971–986, 2000.

[15] J. E. Mathieu, T. S. Heffner, G. F. Goodwin, J. A. Cannon-Bowers, and E. Salas. Scaling the quality of teammates' mental models: equifinality and normative comparisons. *Journal of Organizational Behavior*, 26(1):37–56, 2005.

[16] J. E. Mathieu, T. S. Heffner, G. F. Goodwin, E. Salas, and J. A. Cannon-Bowers. The influence of shared mental models on team process and performance. *Journal of Applied Psychology*, 85(2):273–283, 2000.

[17] J. P. Michael, M. Montemerlo, M. Pollack, N. Roy, and S. Thrun. Towards robotic assistants in nursing homes: Challenges and results. *Robotics and Autonomous Systems*, 42:271–281, 2002.

[18] M. N. Nicolescu and M. J. Mataric. Natural methods for robot task learning: Instructive demonstrations, generalization and practice. In *In Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 241–248, 2003.

[19] J. Pineau and S. Thrun. An integrated approach to hierarchy and abstraction for pomdps. Technical report, 2002.

[20] A. Rowe, N. Cooke, E. Hall, and T. Halgren. Toward an on-line knowledge assessment methodology: Building on the relationship between knowing and doing. *Journal of Experimental Psychology: Applied*, 2:31–47, 1996.

[21] E. Salas, J. Fowlkes, R. Stout, D. Milanovich, and C. Prince. Does crm training improve teamwork skills in the cockpit?: Two evaluation studies. *Human Factors*, 41:326–343, 1999.

[22] J. Shah and C. Breazeal. An empirical analysis of team coordination behaviors and action planning with application to human-robot teaming. *Human Factors*, 52, 2010.

[23] A. Stopp, S. Horstmann, S. Kristensen, and

F. Lohnert. Towards interactive learning for manufacturing assistants. In *In Proc. of the IEEE Int. Workshop on Robot and Human Interactive Communication RO-MAN ï£¡01. Bordeaux/Paris*, pages 18–21, 2001.

[24] R. Stout, J. Cannon-Bowers, E. Salas, and D. Milanovich. Planning, shared mental models, and coordinated performance: an empirical link established. *Human Factors*, 41:61–71, 1999.

[25] S. Szabo, R. Norcross, and W. Shackleford. Safety of human-robot collaboration systems. `http://www.nist.gov/el/isd/ps/safhumrobcollsys.cfm`.

[26] M. van Lent and J. E. Laird. Learning procedural knowledge through observation. In *Proceedings of the 1st international conference on Knowledge capture*, K-CAP '01, pages 179–186, New York, NY, USA, 2001. ACM.

[27] C. Volpe, J. Cannon-Bowers, E. Salas, and P. Spector. The impact of cross training on team functioning. *Human Factors*, 38:87–100, 1996.

[28] F. Wallhoff, J. Blume, A. Bannat, W. Rosel, C. Lenz, and A. Knoll. A skill-based approach towards hybrid assembly. *Advanced Engineering Informatics*, 24(3):329 – 339, August 2010.