AAAI Technical Report FS-12-04
Human Control of Bioinspired Swarms

Human-Inspired Techniques for Human-Machine Team Planning

Julie Shah and **Been Kim** and **Stefanos Nikolaidis**
Massachusetts Institute of Technology
Department of Aeronautics and Astronautics
Computer Science and Artificial Intelligence Laboratory
77 Massachusetts Ave.
Cambridge, MA 02139

Abstract

Robots are increasingly introduced to work in concert with people in high-intensity domains, such as manufacturing, space exploration and hazardous environments. Although there are numerous studies on human teamwork and coordination in these settings, very little prior work exists on applying these models to human-robot interaction. This paper presents results from ongoing work aimed at translating qualitative methods from human factors engineering into computational models that can be applied to human-robot teaming. We describe a statistical approach to learning patterns of strong and weak agreements in human planning meetings that achieves up to 94% prediction accuracy. We also formulate a human-robot interactive planning method that emulates cross-training, a training strategy widely used in human teams. Results from human subject experiments show statistically significant improvements on team fluency metrics, compared to standard reinforcement learning techniques. Results from these two studies support the approach of modeling and applying common practices in human teaming to achieve more effective and fluent human-robot teaming.

Introduction

We envision a future in which collaboration between humans and robots will be indispensable to our work in many domains, ranging from manufacturing to surgery to space exploration. The success of these systems will depend in large part on the ability of robots to integrate with existing human teams. Our goal is to develop robot partners that we can work with more easily and naturally, as inspired by the way we work with other people.

Our hypothesis is that the performance of human-robot teams is improved when a robot teammate emulates the effective teamwork behaviors observed in human teams. There is a precedent for human-human interaction (HHI) research informing the design for human-robot interaction, e.g. (Lockerd and Breazeal 2004; Sakita et al. 2004; Sidner et al. 2005; Trafton et al. 2005; Hoffman and Breazeal 2007). We draw from a body of human-human interaction (HHI) research that has not yet been widely applied to HRI: studies in human teaming for high-intensity domains, including studies of military tactical teams (Entin, Serfaty, and Deckert 1994; Entin and Serfaty 1999; Stout et al. 1999), aviation crews (Salas et al. 1999), medical teams (Mackenzie, Xiao, and Horst 2004), and disaster response crews (Stent 2000). Team dynamics have a significant impact on performance within these domains, producing a strong incentive for teams to understand and apply the communication and coordination strategies that improve performance.

This paper presents recent results from ongoing work aimed at translating qualitative methods from human factors engineering and human team coordination into computational models that can be applied to human-machine team planning. We describe two studies. The first study presents a statistical approach to learning patterns of strong and weak agreements in human planning meetings. Our approach applies statistical machine learning to dialog features, which prior studies in cognitive psychology have shown qualitatively capture the level of commitment to plan choices. Initial results indicate that we can achieve up to 94% average accuracy in predicting strong and weak agreement. This work is the first step towards designing an intelligent robot partner that participates in natural human team planning sessions, encouraging team members to revisit decisions that may adversely affect the team's performance, and spurring dialog that results in higher quality plans.

The second study describes the design and evaluation of computational teaming models that support human-robot cross-training. Cross-training is a technique widely used in human teams, whereby team members iteratively switch roles on the team. This method has been empirically validated to improve mental model similarity among human team members and to improve team performance measures. We formulate human-robot cross-training and evaluate it in a user study $(n = 36)$. Results show statistically significant results on quantitative human-robot mental model elicitation measures and teamwork fluency metrics, compared to standard reinforcement learning techniques where the human assigns rewards to the robot. Additionally, significant differences emerge in subjective measures related to perceived robot performance and human trust. These results support the approach of modeling and applying common practices in human teaming to achieve more effective and fluent human-robot teaming.

Study: Quantitative Prediction of Strength of Agreement in Human Team Planning

Goal-oriented meetings are frequent occurrences in our everyday lives. For example, we discuss project plans at work and coordinate with friends to organize outings. The consequences are inconvenient but relatively minor if the participants leave these types of discussions with different understandings of what was decided upon. Yet, in high-intensity domains such as disaster-response, minor differences in understandings may degrade the team's ability to successfully coordinate and may have serious consequences for people's safety.

It is challenging for a group to reach consensus through dialog. The interaction among people is complicated, dynamic and unpredictable. Natural human collaborative dialog unfolds in cycles, agreements are fluid, and proposals are often implicitly communicated and accepted (Eugenio et al. 1999). In addition, there are social aspects that are often hard to formulate into equations. These characteristics make explicit modeling and quantitative analysis of goal-oriented dialog challenging.

Prior art provides a theoretical foundation for translating the ambiguous and inconsistent nature of dialog into a set of *dialog features* that indicate that strength of agreement to plan decisions (Eugenio et al. 1999). In our recent work, we generalize this qualitative approach to enable a quantitative, predictive capability for characterizing weak and strong agreement in dialog. We envision this work a first step towards the design of an intelligent agent or robot that observes human team planning and interjects to highlight weak agreement among team members. This approach would integrate an intelligence agent into natural human team dynamics and does not require extensive codifying of domain knowledge. In contrast, prior approaches to decision support in planning utilize automated planners to provide suggestions to the team. A common criticism of this approach is that automated planners can not practically capture all relevant domain knowledge and expertise and frequently provide uninformative solutions.

Our approach to predicting strength of agreement builds on prior, qualitative investigations of the human team decision-making process ((Eugenio et al. 1999), (Black 1948), (Hiltz, Johnson, and Turoff 1986)). In particular, Eugenio et al. (Eugenio et al. 1999) present an empirical study of the use of *dialog features* for characterizing strength of agreement. The dialog features capture the level of joint commitment as the negotiation unfolds, from an inability to commit due to lack of full information (*partner decidable option*), to conditional commitment (*proposal*), to unconditional commitment. Each of these dialog features are composed in part of traditional dialog acts, which represent the intention or the role of an utterance in dialogue. Dialogue acts are widely in natural language processing systems for training and testing (Nagata and Morimoto 1994; Stolcke et al. 2000; Ji and Bilmes 2005). Eugenio et al. argue, and we validate quantitatively, that dialogue acts (in particular, accept and reject) are not sufficient for categorizing strength of agreement in dialog. This is because dialog acts only consider one agent's attitude toward an action and therefore do not provide information on the level of joint commitment. In addition, the traditional dialogue acts are not able to capture implicit accepts and rejects in the dialogue.

In our recent work, we utilize both dialogue acts and Eugenio dialog features within a machine learning framework to quantitatively predict strength of agreement. A key benefit to using dialog acts and features to characterize strength of agreement is that this approach does not require the extraction of keywords or other content information from the dialog. In other words, this approach uses information about how the team plans, but does not require storing and processing potentially sensitive information about what they are planning, as is the case for previous quantitative approaches (Hahn, Ladner, and Ostendorf 2006; Hillard and Ostendorf 2003) to this problem. To our knowledge, our approach is the first to (1) estimate strength of agreement based solely on dialog acts and features and not keywords, and to (2) map the qualitative theoretical foundations for strength of agreement to a quantitative, predictive measure.

We apply a number of statistical machine learning techniques to predict strength of agreement, including SVMs, logistic regression, Kmeans, and expectation maximization with gaussian mixture models. We show that we can achieve acceptable accuracy, up to 94% correct prediction of the strength of agreements without using keywords. We also show that Eugenio et al. dialog features play a significant role in prediction accuracy, as compared to analysis using dialog acts alone. We apply SVM feature ranking and Fisher's exact test to analyze the significance of association between classification features (which include Eugenio's features and dialog acts) and strength of agreement. The SVM ranking test indicates Eugenio's features, *Proposal* and *Partner Decidable Option* rank as the top two classification features for the weak agreement category, and Eugenio's feature *Unendorsed Option* ranks among the top two classification features for the strong agreement category. Fisher's Exact Test, a method that is independent of classification technique, indicates that three out of four of Eugenio's features are ranked among the top five classification features. These results lend support to the hypothesis that Eugenio's features play a significant role in the prediction task.

Study: Human Team Training as a Guide for Human-Robot Team Training

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) (Cooke, Salas, and Cannon-bowers 2000) 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 (Mathieu et al. 2005) has shown that teams with accurate but different 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 our recent work, we design and evaluate 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. Our approach to human-robot team training uses an cross-training phase (Stout et al. 1999; Volpe et al. 1996; Marks et al. 2002), which preceeds task execution. There are three types of cross-training (Blickensderfer E. and E. 1998) a) positional clarification (b) positional modeling and c) positional rotation. Findings (Marks et al. 2002; Volpe et al. 1996; Cannon-Bowers J.A. 1998) suggest that positional rotation, which is defined as "learning interpositional information by switching work roles", is the most strongly correlated to improvement in team performance, as it "provides hands on approach to learning interpositional information by giving members experience on carrying out teammates' duties through active participation in each member's role" (Marks et al. 2002). The goal of positional rotation is to provide the individual with hands-on knowledge about the roles and responsibilities of other teammates, with the purpose of improving interrole knowledge and team performance.

We emulate positional rotation in human teams by having the human and robot iteratively switch roles. We name the phase where the roles of the human and robot match the ones of the actual task execution as the *forward phase*, and the phase where human and robot roles are switched as *rotation phase*. In order for the robot's computational teaming model to converge to the human mental model:

1. The robot needs to have an accurate estimate of the human's role in performing the task, and this needs to be similar to the human's awareness of his or her own role. Based on the above, we use the human-robot forward phase of the training process to update our estimation of the transition probabilities that encode the expected human behavior.
2. The robot's actions need to match the expectations of the human. We accomplish this by using the human inputs in the rotation phase to update the reward assignments.

We computationally encode the human-robot teaming model as a Markov Decision Process (MDP) and show that our formulation captures knowledge about the role of the robot and human team member and is quantitatively comparable to the human mental model. Additionally, we propose quantitative measures to assess human-robot mental model convergence, as it emerges through a training process, and mental model similarity. We then introduce a human-robot interactive planning method which uses the MDP computational teaming model to emulate cross-training (Marks et al. 2002).

We compare human-robot cross-training to standard reinforcement learning algorithms through a large-scale user-study of 36 human subjects. Specifically, we compare the proposed formulation to the standard interactive reinforcement learning approach, where the reward signal is provided by a human teacher or coach (Russell and Norvig 2003). In this work, we benchmark against the reinforcement learning algorithm Sarsa(λ) with greedy policy (Sutton and Barto 1998). We chose Sarsa(λ) for its popularity and applicability in a wide variety of tasks. In particular, Sarsa(λ) has been used to benchmark TAMER framework (Knox and Stone 2009), as well as to test TAMER-RL (Knox and Stone 2010; 2012). Furthermore, if we remove eligibility traces (by setting $\lambda = 0$), our implementation of Sarsa(λ) with greedy policy is identical to the Q-Learning with Interactive Rewards (Thomaz and Breazeal 2006). Additionally, variations of Sarsa have been used to teach a mobile robot to deliver objects (Ramachandran and Gupta 2009), for navigation of a humanoid robot (Navarro, Weber, and Wermter 2011), as well as in an interactive learning framework, where the user gives rewards to the robot through verbal commands (Tenorio-Gonzalez, Morales, and Villaseñor Pineda 2010).

Our experiment results show that cross-training improves quantitative measures of human-robot mental model convergence ($p = 0.04$) and mental model similarity ($p = 0.03$), as compared to Sarsa. Additionally, a post-experimental survey shows statistically significant differences in the perceived robot performance, as well as the trust to the robot ($p < 0.01$). Finally, we observed a significant improvement in team fluency metrics, such as an increase of 71% in concurrent motion ($p = 0.02$) and a decrease of 41% in human idle time ($p = 0.04$), during the actual human-robot task execution phase that succeeded the human-robot interactive planning process.

Conclusions

This paper presents results from ongoing work aimed at translating qualitative methods from human factors engineering and human team coordination into computational models that can be applied to human-machine team planning. We describe results from two studies. The first applies statistical machine learning to quantitatively predict weak and strong agreement in human planning meetings. The approach uses statistical analysis of dialog features, which prior studies in cognitive psychology have shown to qualitatively capture the level of commitment to plan choices. Initial results indicate that we can achieve up to 94% average accuracy in predicting strong and weak agreement. The second study designs a computational teaming model and formulates a human-robot interactive planning method that emulates cross-training, a training strategy widely used in human teams. Results from human subject experiments show statistically significant results on team fluency metrics, compared to standard reinforcement learning techniques. These results support the approach of modeling and applying common practices in human teaming to achieve more effective and fluent human-robot teaming.

Acknowledgements

This work is supported in part by MIT Lincoln Laboratory, performed in collaboration with Larry Bush, and by ABB Inc., performed in collaboration with Thomas Fuhlbrigge, Gregory Rossano, Carlos Martinez, and Biao Zhang of ABB Inc., USCRC - Mechatronics. We would also like to acknowledge the Onassis Foundation as a sponsor.

References

Black, D. 1948. On the Rationale of Group Decision-making. *Journal of Political Economy* 56(1):23–34.

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

Cannon-Bowers J.A., Salas E., B. E. B. C. 1998. The impact of cross-training and workload on team functioning: a replication and extension of initial findings. *Human Factors* 92–101.

Cooke, N. J.; Salas, E.; and Cannon-bowers, J. A. 2000. Measuring team knowledge. *Human Factors* 42:151–173.

Entin, E., and Serfaty, D. 1999. Adaptive team coordination. *Human Factors* 41:312–325.

Entin, E.; Serfaty, D.; and Deckert, J. 1994. *Report No TR-648-1*. Burlington, MA: ALPHATECH.

Eugenio, B. D.; Jordan, P. W.; Thomason, R. H.; and Moore, J. D. 1999. The agreement process: An empirical investigation of human-human computer-mediated collaborative dialogues.

Hahn, S.; Ladner, R.; and Ostendorf, M. 2006. Agreement/Disagreement Classification: Exploiting Unlabeled Data using Contrast Classifiers. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, 53–56. New York City, USA: Association for Computational Linguistics.

Hillard, D., and Ostendorf, M. 2003. Detection of agreement vs. disagreement in meetings: Training with unlabeled data. In *In Proc. HLT-NAACL Conference*.

Hiltz, S. R.; Johnson, K.; and Turoff, M. 1986. Experiments in group decision making communication process and outcome in face-to-face versus computerized conferences. *Human Communication Research* 13(2):225–252.

Hoffman, G., and Breazeal, C. 2007. Cost-based anticipatory action-selection for human-robot fluency. *IEEE Transactions on Robotics and Automation* 23(5):952–961.

Ji, G., and Bilmes, J. 2005. Dialog act tagging using graphical models. In *Proc. ICASSP*, volume 1, 33–36.

Knox, W. B., and Stone, P. 2009. Interactively shaping agents via human reinforcement: The tamer framework. In *The Fifth International Conference on Knowledge Capture*.

Knox, W. B., and Stone, P. 2010. Combining manual feedback with subsequent mdp reward signals for reinforcement learning. In *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*.

Knox, W. B., and Stone, P. 2012. Reinforcement learning from simultaneous human and mdp reward. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*.

Lockerd, A., and Breazeal, C. 2004. Tutelage and socially guided robot learning. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3475–3480.

Mackenzie, C.; Xiao, Y.; and Horst, R. 2004. Video task analysis in high performance teams. *Cognition, Technology, and Work* 6:139–147.

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

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

Nagata, M., and Morimoto, T. 1994. First steps towards statistical modeling of dialogue to predict the speech act type of the next utterance. *Speech Communication* 15(3-4):193–203.

Navarro, N.; Weber, C.; and Wermter, S. 2011. Real-world reinforcement learning for autonomous humanoid robot charging in a home environment. In *Proceedings of the 12th Annual conference on Towards autonomous robotic systems*, TAROS'11, 231–240. Berlin, Heidelberg: Springer-Verlag.

Ramachandran, D., and Gupta, R. 2009. Smoothed sarsa: reinforcement learning for robot delivery tasks. In *Proceedings of the 2009 IEEE international conference on Robotics and Automation*, ICRA'09, 3327–3334. Piscataway, NJ, USA: IEEE Press.

Russell, S. J., and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach*. Pearson Education.

Sakita, K.; Ogawara, K.; Murakami, S.; Kawamura, K.; and Ikeuchi, K. 2004. Flexible cooperation between human and robot by interpreting human intention from gaze information. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 846–851.

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

Sidner, C.; Lee, C.; Kidd, C.; Lesh, N.; and Rich, C. 2005. Explorations in engagement for humans and robots. *Artificial Intelligence* 166(1):140–164.

Stent, A. J. 2000. The monroe corpus. Technical report, University of Rochester.

Stolcke, A.; Ries, K.; Coccaro, N.; Shriberg, E.; Bates, R.; Jurafsky, D.; Taylor, P.; Martin, R.; Ess-Dykema, C.; and Meteer, M. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational linguistics* 26(3):339–373.

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

Sutton, R. S., and Barto, A. G. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Tenorio-Gonzalez, A. C.; Morales, E. F.; and Villaseñor Pineda, L. 2010. Dynamic reward shaping: training a robot by voice. In *Proceedings of the 12th Ibero-American conference on Advances in artificial intelligence*, IBERAMIA'10, 483–492. Berlin, Heidelberg: Springer-Verlag.

Thomaz, A. L., and Breazeal, C. 2006. Reinforcement learning with human teachers: evidence of feedback and guidance with implications for learning performance. In *Proceedings of the 21st national conference on Artificial intelligence - Volume 1*, AAAI'06, 1000–1005. AAAI Press.

Trafton, J.; Cassimatis, N.; Bugajska, M.; Brock, D.; Mintz, F.; and Schultz, A. 2005. Enabling effective human-robot interaction using perspective-taking in robots. *IEEE Transactions on Systems, Man, and Cybernetics* 35(4):460–470.

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