

Adapting Plans through Communication with Unknown Teammates

(Doctoral Consortium)

Trevor Santarra
Advisor: Arnav Jhala
University of California Santa Cruz
Santa Cruz, CA 95064
tsarratt@soe.ucsc.edu

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Cooperation and coordination

Keywords

Ad Hoc Teams; Coordination; Multiagent Communication

1. INTRODUCTION

Coordinating a team of autonomous agents is a challenging problem. Agents must act in such a way that makes progress toward the achievement of a goal while avoiding conflict with their teammates. In information asymmetric domains, it is often necessary to share crucial observations in order to collaborate effectively. In traditional multiagent systems literature, these teams of agents share an identical design for reasoning, planning, and executing actions, allowing perfect modeling of teammates. Ad hoc teamwork [10] further complicates this problem by introducing a variety of teammates with which an agent must coordinate. In these scenarios, one or more agents within a team can be unfamiliar, having unknown planning capabilities guiding their behavior.

Much of the existing ad hoc teamwork research focuses on reinforcement learning and decision-theoretic planning. Agents use models of known behavior to predict an ad hoc agent's actions, using decision theory to maximize expected utility in instances where the predicted actions are uncertain. Online learning refines these models with observations of behaviors during execution, increasing the accuracy of the models' predictions, permitting the team to coordinate more effectively [1].

My thesis addresses the problem of planning under teammate behavior uncertainty by introducing the concept of intentional multiagent communication within ad hoc teams. In partially observable multiagent domains, agents much share information regarding aspects of the environment such that uncertainty is reduced across the team, permitting better coordination. Similarly, we consider how communication may be utilized within ad hoc teams to resolve be-

havioral uncertainty. Transmitting intentional messages allows agents to adjust predictions of a teammate's individual course of action. In short, an ad hoc agent coordinating with an unknown teammate can identify uncertainties within its own predictive model of teammate behavior then request the appropriate policy information, allowing the agent to adapt its personal plan. The main contribution of this work is the characterization of the interaction between learning, communication, and planning in ad hoc teams.

2. INFERENCE OVER BEHAVIORS

Prior experience with a variety of agents can provide a set of known behavior models for future use. Observations of an unknown teammate can then be used to identify a similar existing model as a predictor of its future actions. Our initial work in ad hoc teams considered the case where an agent coordinates with a teammate whose behavior is not represented well by a single model but rather by a composition of modeled behaviors. In this scenario, the teammate may adopt new strategies on the fly, adapting its behavior or pursuing new goals unexpectedly. Traditional probabilistic approaches assume static behavior and do not correct the choice of model in a responsive manner [1]. Furthermore, perfect Bayesian inference would require either known transition probabilities, which are not provided in ad hoc domains, or a mechanism for learning the transition likelihoods. By adopting a technique from on-line learning [2], we demonstrated a computationally lightweight technique for responsively identifying changes in a teammate's behavior, correctly selecting an alternate model that most accurately expresses the post-change observed actions. This approach outperforms the standard Bayesian weight revision method, showing promise for accurately modeling agents of inconsistent behavior as well as agents whose high level strategy can be expressed by a sequence of more simple behaviors [6, 8].

Such model inference approaches, while suitable for a wide range of problems, are not without their limitations, however. The possession of a set of known models requires either prior experience or hand-authored knowledge. Even when complex behavior can be approximated with the composition of known, simple models, the low level behaviors must be available prior to coordination. Initial work in learning a model during coordination has relied on generalization of information across known models to speed up the process [1], and as a result, such transfer learning techniques may not apply to novel conditions in which no prior knowledge is

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.

Copyright © 2016, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

available. Furthermore, relying on observations necessitates an act be performed before the information can be utilized to infer or construct an accurate model, a constraint unsuitable for domains with strict limits on potential attempts of completing a task. For this reason, we motivate communicating *in advance* of a point of uncertainty.

3. COMMUNICATION

Across many communicative multiagent frameworks, such as the COM-MTDP model [4] and STEAM [11], communicative actions are often limited to sharing observations. As agents in such systems have complete information regarding the planning capacities of their teammates, they can simply use the shared information to compute precisely how all agents will act. Since the policies of teammates are the source of uncertainty in ad hoc teams, it follows that policy information is a promising target for communicative acts. We introduced this idea in [7].

In early decision-theoretic agent communication literature, various types of communication were theoretically validated in their effect on coordinating multiple agents. These included intentional messages, questions, proposals, threats, imperatives, and statements of propositional attitudes [3]. In each case, providing or requesting information adjusted one or more agents' policies through refining an agent's expectations of either its own policy's likelihood of success or the intentions of another agent acting within the same environment. Analogously, the refinement of predicted action probabilities and, consequently, an improved policy for a coordinating agent is desirable for ad hoc teams.

Whereas the broadcast of the intention of pursuing a goal addresses multiple state-action pairs within an agent's policy computation, we must consider that an unfamiliar teammate may not possess the capability of reasoning with high level abstractions such as joint plans or hierarchical goal decompositions. However, we put forth the observation that all agents involved are universally tasked with assigning actions to states, independent of the particular planning implementation details. From a general perspective, we consider how an ad hoc agent could benefit from obtaining a single state-action pair—the atomic component of a teammate model—from communication.

In our extended abstract to be presented at *AAMAS-16* [9], we have provided an efficient procedure for evaluating potential state-action pairs for communication by examining the uncertainty within a teammate model as well as empirically analyzed various aspects of the communicative capability. For our test domain, we used a variation of the multiagent pursuit domain where two agents attempt to capture a prey within a maze. We were able to show the trade-off between collected information and communication rates and, likewise, the effect of communication costs on query rates and the resulting expected utility of the agent. Finally, our analysis determined that branch points in a maze are commonly communicated more frequently than neighboring cells, indicating that the tested agents determined resolving uncertainty at such states was associated with higher utility.

4. FUTURE WORK

While our previous work has focused on the ability to coordinate with computer agents, we intend to test our approach with human teammates, given the potential applicability of

ad hoc team approaches to human-computer teams. Furthermore, over the ensuing year, we intend to transition to a more complex domain, using a slightly simplified form of the real-time strategy game, *StarCraft*.

An immediate extension to this work is the consideration of communicating multiple state-action pairs without independent evaluation. It is possible for two states to have no utility for communication individually but have non-zero utility when considered together. This opens up a combinatorial space of potential intentional information sets that could be communicated, similar to problem of picking a subset of observations to share within a team, as explored by Roth et al. [5]. Due to the intractable nature of the problem, the authors motivated the exploration of heuristics as approximate solutions. It is yet unclear how states in such collections will be related, though we hypothesize they could take any one of various forms, including sequences of successive states (forming a plan), areas of connected states (a local subspace of the domain), or groups of independently valuable states (such as branch points discussed earlier).

REFERENCES

- [1] S. Barrett, P. Stone, S. Kraus, and A. Rosenfeld. Teamwork with limited knowledge of teammates. In *AAAI*, 2013.
- [2] O. Bousquet and M. K. Warmuth. Tracking a small set of experts by mixing past posteriors. *The Journal of Machine Learning Research*, 3:363–396, 2003.
- [3] P. J. Gmytrasiewicz, E. H. Durfee, and J. Rosenschein. Toward rational communicative behavior. In *AAAI Fall Symposium on Embodied Language*, pages 35–43, 1995.
- [4] D. V. Pynadath and M. Tambe. The communicative multiagent team decision problem: Analyzing teamwork theories and models. *Journal of Artificial Intelligence Research*, pages 389–423, 2002.
- [5] M. Roth, R. Simmons, and M. Veloso. What to communicate? execution-time decision in multi-agent pomdps. In *Distributed Autonomous Robotic Systems 7*, pages 177–186. Springer, 2006.
- [6] T. Sarratt and A. Jhala. Rapid: A belief convergence strategy for collaborating with inconsistent agents. In *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [7] T. Sarratt and A. Jhala. The role of models and communication in the ad hoc multiagent team decision problem. In *Proceedings of the Third Annual Conference on Advances in Cognitive Systems Poster Collection*, page 27, 2015.
- [8] T. Sarratt and A. Jhala. Tuning belief revision for coordination with inconsistent teammates. In *AIIDE*, 2015.
- [9] T. Sarratt and A. Jhala. Communicating intentions for coordination with unknown teammates. In *Workshops at the Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [10] P. Stone, G. A. Kaminka, S. Kraus, J. S. Rosenschein, et al. Ad hoc autonomous agent teams: Collaboration without pre-coordination. In *AAAI*, 2010.
- [11] M. Tambe. Agent architectures for flexible. In *Proc. of the 14th National Conf. on AI, USA: AAAI press*, pages 22–28, 1997.