Considering Agent and Task Openness in Ad Hoc Team Formation (Extended Abstract)

Bin Chen, Xi Chen, Anish Timsina and Leen-Kiat Soh
Department of Computer Science and Engineering
University of Nebraska, Lincoln, NE
{bchen, xgchen, atimsina, lksoh}@cse.unl.edu

ABSTRACT

When deciding which ad hoc team to join, agents are often required to consider rewards from accomplishing tasks as well as potential benefits from learning when working with others, when solving tasks. We argue that, in order to decide when to learn or when to solve task or both, agents have to consider the existing agents' capabilities and tasks available in the environment, and thus agents have to consider *agent and task openness*—the rate of new, previously unknown agents (and tasks) that are introduced into the environment. We further assume that agents evolve their capabilities intrinsically through learning by observation or learning by doing when working in a team. Thus, an agent will need to consider which task to do or which team to join would provide the best situation for such learning to occur. In this paper, we develop an auction-based multiagent simulation framework, a mechanism to simulate openness in our environment, and conduct comprehensive experiments. Our results, based on more than 20,000 simulation runs, show that considering environmental openness is beneficial and necessary, and task selection strategies leveraging openness can improve agent learning and performance. We also report on observations of emergent behaviors related to openness.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence – *intelligent agents, multiagent systems*

General Terms

Algorithms, Measurement, Performance, Design, Experimentation

Keywords

Ad hoc Team Formation; Agent Openness; Task Openness; Learning by Doing; Learning by Observation

1. INTRODUCTION

Many aspects of ad hoc team formation have been studied, focusing on learning, leading, and dealing with uncertainties in agent behavior [1-4]. But, as we try to study team formation in certain agents, like human, we need to consider several factors like how human learn from working in a team as well as observing a teammate. Research done so far, while considering learning [4], has not considered the learning that is present when agents—such as humans—work together in a team. For example, when human agents work together, it is inevitable that they learn from each other. In ad hoc team formation, while prior knowledge of a potential teammate is not available, it is still possible for an agent to model the types of agents and tasks likely to be in the environment, and to assume that learning is inevitable when working together. Such consideration and assumption will influence how agents form ad hoc teams—in how each decides to join an ad hoc team to help solve a task.

Furthermore, another key question is how agents should decide on which teams to join when taking into account the potential rewards of learning while on a team. An agent would have to tradeoff between combined reward resulting from optimizing on task rewards and that resulting from optimizing on learning. In an ad hoc environment where an agent has little or no knowledge about potential teammate, how should such an agent leverage what it can model of the environment to help make this decision?

We see that there are two types of openness from a multiagent viewpoint [3]. First, *task openness* (*TO*) refers to the rate of new, previously unseen tasks that are introduced into the environment. Second, *agent openness* (*AO*) refers to the rate of new, previously unknown agents that are introduced into the environment, while known agents exit the environment. For example, an agent whose particular capability is low may choose to join a team with a good opportunity to learn about this capability from other teammates (via observation) even when the direct rewards of completing this task is low. Thus if the degree of agent openness is high, then the likelihood to work with the same agent/agent type to learn about a particular capability would be low. So it might be prudent for the agent to lean towards joining a team to learn from the particular agent/agent type sooner than later. Or if agent openness is low, an agent might want to become an expert at a particular capability by repeatedly completing a particular task to learn by doing such that it can leverage other agents' expertise to collaboratively solve tasks more effectively. On the other hand, if task openness is high, such that different tasks appear and disappear from the environment very often, then the likelihood of encountering the same task/task type again would be low, then agents do not have to spend time, effort, and resource to learn to solve a particular task/task type—say, a difficult one—if the task/task type would not likely appear again in the future. In that case, an agent might not care too much about learning to solve that task/task type, and instead aim for getting more direct rewards and sooner.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, Weiss, Yolum, Bordini, and Elkind (eds.), May, 4–8, 2015, Istanbul, Turkey.

2. TASK SELECTION STRATEGIES

In our simulation design, tasks are allocated through auctions held on blackboard. Agents can see the available tasks as well as tasks' specification. Then, based on this information and agents' perception of AO and TO, agents decide which tasks to bid on. The following task selection strategies are based on the assumption that the system administrator assigns each subtask $\tau \in \mathcal{T}$ to the best qualified agents who bid on the task $\mathcal{T}$.

Strategy 1. Most Qualified (MQ). This strategy finds the task that gives the maximum total positive differences of agent a_i's corresponding capabilities of subtasks and the quality requirement of subtasks in each task $\mathcal{T}$.

Strategy 2. Most Learning Opportunity (MLO). This strategy finds task with best potential utility that the bidding agent can gain by observing other teammates solving subtasks.

Strategy 3. Most Qualified + Learning (MQ+LO). This strategy is a hybrid of the first two strategies. Agents consider the opportunity to learn from other agents by observation and their qualification for solving one subtask within a task.

Strategy 4. Most Total Potential Utility (MTPU).

$$U(\mathcal{T}) = w_L \cdot U_{learn}(\mathcal{T}) + w_S \cdot U_{solve}(\mathcal{T}) \quad (1)$$

where w_L and w_S are the weights for learning and solving a task, respectively, and $w_L + w_S = 1$. $U_{learn}(\mathcal{T})$ is the potential utility from learning by doing and learning by observation as used in ***MQ+LO***. $U_{solve}(\mathcal{T})$ is the potential utility of the bidding agent participating in solving the task $\mathcal{T}$, This strategy finds the task that gives the maximum of $U(\mathcal{T})$. We have several interesting variants by setting the weights differently: ***Strategy 4.1. MTPU_L=S*** with $w_L = w_S = 0.5$; ***Strategy 4.2. MTPU_L<S*** with $w_L = 0.25, w_S = 0.75$, and ***Strategy 4.3. MTPU_L>S*** with $w_L = 0.75, w_S = 0.25$;

Strategy 5. MTPU with Agent Openness (MTPU+AO). This strategy is based on Eq.1, but taking AO into account. Hence, for the *MTPU+AO* strategy, we set $w_L = AO$ and $w_S = 1 - AO$.

Strategy 6. MTPU with Task Openness (MTPU+TO). Similarly, using the same Eq. 1, but taking TO into account. When TO is high, focusing on immediate rewards is a good choice. Hence for this strategy, we set $w_L = 1 - TO$ and $w_S = TO$.

Strategy 7. MTPU with Both Openness (MTPU+ATO). Similarly, using the same Eq. 1 for the *MTPU+ATO* strategy, we use $w_L = \frac{AO}{AO+TO}$ and $w_S = \frac{TO}{AO+TO}$. We define w_L and w_S in Strategy 7 as such so that when AO and TO are either high or both low, the weight for learning (w_L) and the weight for getting the immediate rewards (w_S) are not too different from each other.

3. RESULTS

We conducted three experiments to investigate (1) the impact of openness in task completion and learning, (2) the performance of different task selection strategies in open environment and (3) the impact of initial agent expertise in task completion and learning.

Our simulation results shows that it becomes difficult for agents to complete tasks or learn if the environment is open. High TO caused new tasks to be introduced to the system, which required skills that were more varied. In such case, higher AO generally helped agents complete more task and lower AO did the opposite.

Moreover, we observed that strategies that considered AO and TO performed much better than those that did not. As these strategies leveraged the agent's model of the environmental openness, they could make the decision to learn or go for immediate rewards more effectively. We observed that when both agents and tasks were more open, strategy *MTPU_L<S (75% Solve, 25% Learn)* performed the best. This is because when AO is non-zero, new agents were introduced and old agents left with their learned capabilities. As a result, agents had fewer opportunities to use their new learned capabilities to solve tasks before they left. Moreover, high TO meant learned capabilities might not be used as tasks requiring newer capabilities were introduced in the system.

Lastly, we observed that if agents had just 1 non-zero initial capabilities—i.e., agents were not very capable, then strategy *MLO* performed the best as it gave agents the chance to learn more capabilities. On the other hand, when agents had 5 non-zero initial capabilities—i.e., agents were very capable, then all strategies which focused on solving performed better. This was because in the former case, agents among themselves did not have enough initial capabilities to solve all tasks efficiently –hence the need to learn, whereas in the latter case agents already had enough initial capabilities to solve tasks –hence they focus on solving tasks.

4. CONCLUSIONS

We have developed an ad hoc team formation framework that takes into account learning and task solving under varying degrees of environmental openness. Running simulations of this framework, we were able to study various effects of considering agent openness (*AO*) and task openness (*TO*) in ad-hoc team formation. We observed that *AO* and *TO* are important in ad hoc team formation. First, we saw that, *AO* and *TO* change the way teams are formed. In open environment, agents need to factor in the possibility of new agents and tasks entering the environment in order to make better decisions in terms of joining a team. Second, *AO* impacts learning, with the introduction of new agents boosting the learning when new tasks are also being introduced. *TO* makes it difficult for agents to solve the tasks. The possibility of new tasks emerging in the environment means newer agents entering the environment can be helpful as they could bring newer capabilities. Having now established the importance of *AO* and *TO*, gaining insights into the relationship between them, and investigating the effectiveness of several openness-based task selection strategies, we have identified several line of research. First, we will *explore more realistic ways to perceive openness*—as our current assumption where agents know both *AO* and *TO* exactly, is not ideal. Second, we will consider the impact of *both* teaching and learning while modeling agent's behavior, particularly incorporating the fundamental game-theoretic work from Stone and Kraus [4].

ACKNOWLEDGMENTS

This research was supported in part by an NSF grant DUE-1431874.

REFERENCES

[1] Barrett, S., P. Stone, S. Kraus, and A. Rosenfeld. Learning teammate models for ad hoc teamwork. In *AAMAS Adaptive Learning Agents Workshop*. 2012.

[2] Barrett, S., P. Stone, and S. Kraus. Empirical Evaluation of Ad Hoc Teamwork in the Pursuit Domain. *AAMAS'2011*, 567–574.

[3] Jumadinova, J., P. Dasgupta, and L.-K. Soh. Strategic Capability-Learning for Improved Multiagent Collaboration in Ad Hoc Environments. *IEEE Trans. Systems, Man, & Cybernetics: Systems,* vol. 44, no. 8 (2014): 1003-1014.

[4] Stone, P., and S. Kraus. To teach or not to teach?: decision making under uncertainty in ad hoc teams. *AAMAS'2010*, 117-124.