Role Selection in Ad Hoc Teamwork

(Extended Abstract)

Katie Genter, Noa Agmon, and Peter Stone
Department of Computer Science
The University of Texas at Austin
{katie,agmon,pstone}@cs.utexas.edu

ABSTRACT

An ad hoc team setting is one in which teammates must work together to obtain a common goal, but without any prior agreement regarding how to work together. In this work we introduce a *role-based approach* for ad hoc teamwork, in which each teammate is inferred to be following a specialized role that accomplishes a specific task or exhibits a particular behavior. In such cases, the role an ad hoc agent should select depends both on its own capabilities and on the roles currently selected by other team members. We present methods for evaluating the influence of the ad hoc agent's role selection on the team's utility and we examine empirically how to choose the best suited method for role assignment in a complex environment. Finally, we show that an appropriate assignment method can be determined from a limited amount of data and used successfully in new tasks that the team has not encountered before.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Algorithms, Experimentation

Keywords

Ad Hoc Teamwork, Agent Cooperation, Coordination

1. INTRODUCTION

Ad hoc teamwork is a relatively new research area [1, 4, 5] that examines how an agent ought to act when placed on a team with other agents such that there was no prior opportunity to coordinate behaviors. This is in contrast to most prior multiagent teamwork research, which often requires explicit coordination protocols, languages, and/or shared assumptions (e.g. [3, 6]).

In some team domains, the team behavior can be broken down into *roles*. In such domains, an ad hoc agent's main task is to decide which role to assume, such that the team's performance is maximized. The decision of which role an ad hoc agent should assume is situation-specific: it depends on the task the team performs, the environment in which it operates, and the capabilities of the team members. One trivial approach to the problem is for an ad hoc agent to assume the role at which it is most *individually* capable. However, the choice of optimal role—one that results in highest *team* utility—rarely depends only on the ad hoc agent, but also on the ability and behavior of the other team members. We therefore examine the contribution of an ad hoc agent to the team by the measure of *marginal utility*, which is the increase in a team's utility when an ad hoc agent is added to the team and assumes a particular role. An *optimal mapping* of an ad hoc agent to a role is, therefore, one that maximizes the marginal utility, hence maximizing the contribution of the ad hoc agent to the team's utility.

2. PROBLEM DEFINITION

An ad hoc teamwork problem is one in which several agents find themselves in a situation where they all have perfectly aligned goals, yet they have had no previous opportunity to coordinate their teamwork [5]. In this work we introduce the *role-based* ad hoc teamwork problem, which is one that requires or benefits from dividing the task at hand into roles. Throughout this paper we refer to the agents that make up a team as either *ad hoc agents* or *teammates*. Ad hoc agents are agents whose behavior we can control, while teammates are agents that we have no control over, potentially because they were programmed at a time when future collaboration with our agents was unforeseeable.

Formally, let task d have m roles $R(d) = \{r_0, ..., r_{m-1}\}$. Let $\mathbf{A} = \{a_0, ..., a_{n-1}\}$ be the set of ad hoc agents and $\mathbf{B} = \{b_0, ..., b_{k-1}\}$ be the set of teammates such that $T = A \cup B$ is the team that is to perform task d. Let mapping $\mathbf{P} : B \rightarrow R(d)$ be the mapping of B to roles $\{r_0, ..., r_{m-1}\}$ and let mapping $\mathbf{S} : A \rightarrow R(d)$ be the mapping of A to roles $\{r_0, ..., r_{m-1}\}$. Finally, let mapping $\mathbf{SP} : T \rightarrow R(d)$ be the combination of mappings S and P. A team score $U(SP, d, T)$ results when the set of agents T perform a task d, with each $t_j \in T$ fulfilling some role $r_i \in R(d)$ under mapping SP. The marginal utility $MU(S, P)$ obtained by mapping S, assuming P is the mapping of B to roles, is the score improvement obtained when S maps A to roles. Hence, marginal utility $MU(S, P) = U(SP, d, T) - U(P, d, B)$.

Given that mapping P is fixed, the role-based ad hoc team problem is to find a mapping S that maximizes marginal utility. Although for the remainder of this paper we focus on the case where $A = \{a_0\}$, the problem definition provided above is valid for any number of ad hoc team agents.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

3. MODELS FOR CHOOSING A ROLE

The gold standard way for an ad hoc agent to determine the marginal utility of selecting a particular role is to determine $U(SP, d, T)$ for each possible role it could adopt. However, in practice, the ad hoc agent must *predict* its marginal utility for all possible roles and then select just *one* role to adopt. Here we lay out three possible models with which the ad hoc agent could do this prediction.

Unlimited Role Mapping Model The value received by the team for an agent performing a role is not dependant on the roles fulfilled by other teammates.

Limited Role Mapping Model The benefit the team recieves for an agent performing role r_i is dependent on the number of agents performing r_i. The team recieves no benefit for an additional agent performing r_i if this results in less than (greater than) r_i^{min} (r_i^{max}) agents performing r_i.

Incremental Role Mapping Model The value added by an agent performing a role is correlated with the number of agents performing that role via a (1) logarithmic, (2) exponential, or (3) sigmoidal function.

4. MODEL EVALUATION

We examine each of the three models described above in a capture-the-flag style variant of Pacman [2]. The Pacman map is divided into two halves and two teams compete by attempting to eat the food on the opponent's side of the map while defending the food on their side. A team wins by eating all but two of the food pellets on the opponent's side or by eating more pellets than the opponent before time expires. The result of each game is a *score differential*—the difference between the number of pellets protected by the team and the number of pellets protected by the opponent.

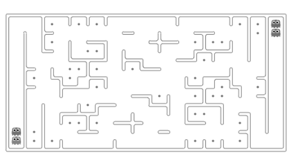

Figure 1: Sample Pacman capture-the-flag map.

4.1 Determining the Best-Suited Model

We use three tasks to determine which of the models best represents the marginal utility of a role selection for the Pacman Capture-the-Flag environment, where a *task* is defined by the number of opponents and the map. In each task we consider two roles that could be performed: $R = \{$offense, defense$\}$.

We start by gathering full sets of *gold standard data*. In particular, we gather score differentials over one thousand games for each team of zero to six offensive agents and zero to six defensive agents (49 teams). In order to emphasize differences in score differentials close to zero, we input the score differential from each game into the sigmoid function $1/1 + e^{-0.13 * scoreDifferential}$ and average the results to obtain gold standard data. Then we use the gold standard data to determine the *gold standard decision* of whether an ad hoc agent should perform an offensive role or a defensive role on any team composed of zero to five offensive agents and zero to five defensive agents. To determine the gold standard decision we look at whether the gold standard data is greater for the team with one extra defensive player or the team with one extra offensive player.

For each of the model functions, we input the gold standard data and the model function into a least squares curve fitting algorithm and obtain *fitted parameters* for the model function. We then use the fitted parameters to calculate *fitted results* for all 49 teams. Lastly, we translate these fitted results into *fitted decisions* using the same methodology used to translate the gold standard data into gold standard decisions. Then we compare the number of times the gold standard decision does not match the fitted decision for a particular team arrangement —in other words, the number of *incorrect decisions*. Our experiments showed that the exponential and sigmoidal functions of the incremental model made the fewest incorrect decisions across the three tasks. Hence we conclude that in the Pacman Capture-the-Flag domain, at least on the maps and opponents studied, the incremental model using an exponential or sigmoidal function most accurately models team utility.

4.2 Predictive Modeling

Once a model type has been selected for a domain, the ad hoc agent can use this model to *predict* the marginal utility of role selection on new tasks for which we have limited gold standard data. We do this by choosing fitted parameters for the new task based on available data. Our experiments showed that the prediction accuracy of each incremental model function variation improved as more data was available, and that some variations did surprisingly well even when provided extremely sparse data.

5. FUTURE WORK

This research is among the first to study role-based ad hoc teams. As such, there are many potential directions for future work. We plan on expanding our work into more complicated environments with more than two potential roles to fulfill and more than one ad hoc agent. Additionally, we wish to consider the case in which the ad hoc agents encounter teammates that are running unfamiliar behaviors, forcing the ad hoc agents to model their teammates.

6. ACKNOWLEDGEMENTS

This work has taken place in the Learning Agents Research Group (LARG) at UT Austin. LARG research is supported in part by NSF (IIS-0917122), ONR (N00014-09-1-0658), and the FHWA (DTFH61-07-H-00030).

7. REFERENCES

[1] M. Bowling and P. McCracken. Coordination and adaptation in impromptu teams. In *AAAI'05*, pages 53–58, 2005.

[2] J. DeNero and D. Klein. Teaching introductory artificial intelligence with pac-man. In *Proceedings of the Symposium on Educational Advances in Artificial Intelligence*, 2010.

[3] B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *AIJ*, 86(2):269 – 357, 1996.

[4] E. Jones, B. Browning, M. B. Dias, B. Argall, M. M. Veloso, and A. T. Stentz. Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks. In *ICRA'06*, pages 570 – 575, 2006.

[5] P. Stone, G. A. Kaminka, S. Kraus, and J. S. Rosenschein. Ad hoc autonomous agent teams: Collaboration without pre-coordination. In *AAAI'10*, 2010.

[6] M. Tambe. Towards flexible teamwork. *JAIR*, 7:83–124, 1997.