Scalable Agent Modeling for Large Multiagent Systems

Carrie Rebhuhn
Oregon State University
rebhuhnc@onid.orst.edu

Introduction

In a heterogeneous multiagent system it can be useful to have knowledge about the different types of agents in the system. Agent modeling develops agent models based on interactions between agents, then predicts agent actions. This approach is effective in small domains but does not scale well. We develop an approach where an agent can learn using an abstract model identification or *stereotype* rather than an explicit and unique model for each agent. We associate each agent with a stereotype and learn a policy incorporating this knowledge. The benefits of this approach are that it is simple, scalable, and degrades gracefully with misidentification.

Research Questions

We have three main goals in this thesis. We first wish to see if, by identifying types in the state space, agents can use this information to increase performance during learning. We then wish to identify system characteristics which recommend the level of agent modeling. Finally we wish to use stereotype identifications with dynamic agents. We do this by asking three research questions:

- **RQ 1**: Can we increase performance in a large multiagent system using type identification and external learning?
- **RQ 2**: Can we identify some aspect of the domain that will indicate the appropriate descriptiveness of the model?
- **RQ 3**: Can we assign types dynamically to account for unconverged learning solutions?

Related Work

Stereotyping has varying definitions in literature. Broadly defined, stereotyping groups several agents under a single model definition. By identifying another agent's policy type, an agent can make a more informed decision. Burnett, Norman, and Sycara combined stereotyping with the concept of trust using a tree model to represent a stereotype (2010). The agent community was then able to share and update these stereotypes. *Stereotypical reputation* allows agents without prior knowledge of others to access group knowledge (Burnett, Norman, and Sycara 2013).

Few stereotype methods explicitly handle dynamic agents, but some accommodate a robust stereotype definition. Denzinger and Hamdan handle potential stereotype misidentification by periodically reevaluating an agent's assigned stereotype (2004). Bard and Bowling implicitly modeling by learning responses to several player types offline. This does not rely on an explicit model of an agent's actions, but does not accommodate dynamic agents (2007).

The novelty of this work is to develop an approach that works with a larger number of adaptive agents than previous approaches.

Proposed Research Plan

We want investigate and develop stereotype methods for large multiagent systems. Toward this, we will show the efficacy of using stereotypes in a large domain. We will then use domain characteristics to define stereotypes. Finally we will adjust our stereotype approach to better accommodate dynamic agents.

Stereotype Use in a Complex Domain

We first show that stereotypes can improve performance in a complex cooperative system: the Conflict-Avoidance domain for Unmanned Aerial Systems (UAS). This problem requires planes to select conflict-avoidance maneuvers while flying continuously between randomized points. The high-level goal is to ensure plane separation while adding minimal extra distance to the plane's path. Formally we define the global utility as:

$$G(\hat{s}, \hat{a}) = \sum_{i \in \mathcal{A}} d_{extra}(a_i) + \frac{|\mathcal{A}|L}{\sum_{i \in \mathcal{A}} d_{neighbor}(s_i)} \qquad (1)$$

where $G(\hat{s}, \hat{a})$ is the global utility dependent on the states $\hat{s}$ and actions $\hat{a}$, $d_{extra}(a_i)$ is the distance cost incurred by the action of agent i from the set $\hat{a}$, $d_{neighbor(s_i)}$ is the euclidean distance of the nearest neighbor of agent i, $|\mathcal{A}|$ is the number of agents, and L is the edge length of the testing map.

We then assess stereotyping against other techniques that focus on noise reduction by comparing to a difference utility. The difference utility reduces noise by comparing the system reward with agent i to a system reward without agent i. This must be accomplished through approximation or resimulation. The performance shown in Figure 1 indicates that the inclusion of stereotypes is more effective than the difference utility at noise reduction in some cases. Using stereotypes requires no resimulation or approximations, which can become computationally expensive. The full results and details for the UAS domain are available in the workshop paper that appeared in ALA 2014 (Rebhuhn, Knudson, and Tumer 2014).

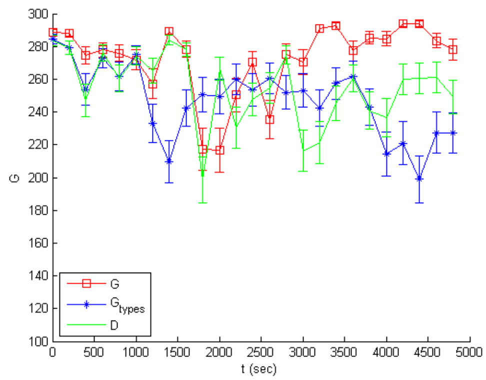

Figure 1: UAS domain comparison of minimizing global reward, global reward with stereotypes, and difference reward.

Developing Stereotype Metrics

In experiments with the UAS domain as well as smaller toy experiments we have identified an important aspect of the application of stereotypes; the *stereotype information value*. If a stereotype does not add important information, it can hinder the learning process by limiting the number of learning examples available. Differentiating between two agent types may not be worth the cost of dividing the number of learning examples. Conversely, making stereotype divisions too coarse loses modeling gains because agent structure becomes environmental noise.

No prescriptive method currently exists for how to define stereotypes to maximize the information of the stereotype. We have already seen the effect of different levels of stereotype differentiation in toy experiments. From these experiments we will 1) define a metric that evaluates the information value given by the stereotype currently in use and 2) use this metric to make appropriate stereotype divisions in a variety of problems. Work toward this should be completed or in progress by January 2015.

Dynamic Stereotype Regeneration

Agents can use stereotypes to categorize fixed policies. An agent that recognizes the pattern of another agent's behavior can take an actions to accommodate this behavior. But large multiagent systems often do not feature agents with fixed policies. Because of this, a stereotyping method for a large multiagent system must accommodate learning agents as well as fixed or converged policies.

We have already included a 'learner' stereotype in the UAS domain, which has an adapting policy. Yet, we do not take advantage of patterns in learning throughout the run, and we do not accommodate the changing demographic of learners. Previous work has explored changes in stereotype assignment, but this work re-estimates an agent's assignment to a set of *fixed* stereotypes (Denzinger and Hamdan 2004). In a large learning system it may be more useful to adapt the stereotype *definition* to remove stereotypes made obsolete by the learning process.

Exploring the adaptation of stereotypes throughout the learning run is important future work. In this work we will focus on how to learn early on with stereotypes of unconverged agents. We will also explore how to incorporate changing system demographics into the learning problem. This will extend the applicability of stereotypes to more dynamic systems.

Conclusion

We show that we can gain advantages from the application of stereotyping in cooperative domains. Future work remains to identify a way to determine a model's value to a system, and to develop methods of handling changing models. With these questions answered, we can construct a more accessible and structured application of stereotypes.

References

Bard, N., and Bowling, M. H. 2007. Particle filtering for dynamic agent modelling in simplified poker. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22-26, 2007, Vancouver, British Columbia, Canada*, 515–521. AAAI Press.

Burnett, C.; Norman, T. J.; and Sycara, K. 2010. Bootstrapping trust evaluations through stereotypes. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, 241–248. International Foundation for Autonomous Agents and Multiagent Systems.

Burnett, C.; Norman, T. J.; and Sycara, K. 2013. Stereotypical trust and bias in dynamic multiagent systems. *ACM Transactions on Intelligent Systems and Technology (TIST)* 4(2):26.

Denzinger, J., and Hamdan, J. 2004. Improving modeling of other agents using stereotypes and compactification of observations. In *3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), 19-23 August 2004, New York, NY, USA*, 1414–1415. IEEE Computer Society.

Rebhuhn, C.; Knudson, M.; and Tumer, K. 2014. Announced strategy types in multiagent rl for conflict-avoidance in the national airspace. In *AAMAS-2014 Workshop on Adaptive and Learning Agents*.