Pragmatic Multi-Agent Learning

Andrew Garland
Computer Science Department
Brandeis University
Waltham, MA 02254
aeg@cs.brandeis.edu

Early models of procedural learning assumed actors were isolated, model-based thinkers. More recently, learning techniques have become more sophisticated as this assumption has been replaced with more realistic ones. To date, however, there has been no thorough investigation of multiple, heterogeneous, situated agents who learn from the pragmatics of their domain rather than from a model. This research focuses on this important problem and develops learning techniques that allow agents to improve their performance in a dynamic environment by learning from past run-time behavior.

Humans provide a natural model of pragmatic agents situated in a multi-agent world. [1] argues that the development of distributed cooperative behavior in people is shaped by the accumulated cultural-historical knowledge of the community. Our learning techniques are motivated by this argument and use a structure called *collective memory* to store the accumulated procedural knowledge of a community of agents. Collective memory contains the breadth of knowledge the community acquires through interacting with each other and the world during the course of solving sequences of distinct problems. The cornerstone of collective memory is a cooperative procedures case-base that augments the agents' first-order planner [2]; in other words, this work follows in the tradition of second-order planners, extending them into multi-agent domains.

In our model of activity, each agent has her own point of view on how best to proceed, which often leads to uncoordinated and unproductive behavior. Furthermore, inefficient behavior would occur even if there was a consensus upon the best course of action to follow (perhaps legislated by a supervising agent or agreed to during community-wide communication) because of the community's initial lack of knowledge about their uncertain domain. Through the use of collective memory, however, agents behave more efficiently over the course of solving a problem sequence for two reasons. First, individual agents develop a point of view based upon shared experiences; second, they learn procedures that capture regularities both in the task environment and in the patterns of cooperation for solving problems in task domain. That is, an agent remembers successful cooperative behavior in which she was involved, and uses it as a basis for future interactions.

In addition to a case-base of cooperative procedures, collective memory currently contains a set of tree structures, called operator probabilities trees [3]. An agent uses these trees to construct higher quality plans by more accurately estimating the probability of success for operators she may attempt. More accurate estimates lead to plans that are more likely to succeed because the estimates are used to guide the first-order planner's search, including the selection of role bindings. Empirical results show that both cooperative procedures and operator probabilities trees lead to significant reductions in the amount of time a community takes to solve randomly generated problems. Furthermore, the two components of collective memory are more effective together than either alone, showing that they facilitate non-overlapping aspects of learning.

Presently, collective memory is implicitly represented as the (conceptual) union of the distributed, private memories of each of the agents in the community. We are developing alternate techniques that would store the procedural knowledge case-base in a central memory (one memory for all agents) or a guild memory (one memory for each group of homogeneous agents). We will conduct experiments to quantify the utility of both of these implementations, as well as of hybrid combinations that incorporate distributed memories.

We believe this research will clearly demonstrate that collective memory is an effective resource for pragmatic multi-agent learning. Hence, it will make a significant contribution to current literature on planning, memory, and group activity.

This work has been advised by Richard Alterman and supported in part by ONR grants N00014-96-1-0440 and N00014-97-1-0604.

References

[1] Michael Cole and Yrjö Engeström. A cultural-historical approach to distributed cognitition. In Gavriel Salomon, editor, *Distributed Cognitions*, pages 1–46. Cambridge University Press, 1993.

[2] Andrew Garland and Richard Alterman. Preparation of multi-agent knowledge for reuse. In *1995 AAAI Fall Symposium on Adaptation of Knowledge for Reuse*, pages 26–33, 1995.

[3] Andrew Garland and Richard Alterman. Multiagent learning through collective memory. In *1996 AAAI Spring Symposium on Adaptation, Coevolution and Learning in Multiagent Systems*, pages 33–38, 1996.