

Learning Team Behaviors with Adaptive Heterogeneity

Liviu Panait
Department of Computer Science
George Mason University
lpanait@cs.gmu.edu

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems

General Terms

Experimentation

Keywords

multiagent learning, dynamic role allocation

The task of hard-coding agent behaviors to achieve desired team behaviors is very difficult, if not intractable. The complexity of multiagent problems can also rise exponentially with the number of agents and their behavioral sophistication. The field of cooperative multiagent learning promises solutions to these issues by employing automatic search methods to discover agent behaviors, and as such it has been the focus of numerous studies in recent years.

Research in cooperative multiagent learning usually concentrates on two opposite approaches, each generating a different range of questions. First, there are techniques to learn homogeneous behaviors for large swarms of identical agents; they assume all agents have identical behavior (a significant reduction to the search space), but the potential of specialization to different tasks is minimized. Second, there are techniques that allow each agent to specialize to a unique behavior, thus creating a wide range of potential solutions; these methods are however applicable to only small teams due to the increased complexity that comes with larger numbers of agents.

My interests target the middle-ground, namely moderately large teams of agents with heterogeneous behaviors. Combining the desideratum for scalability of learning to larger teams with the need for agent specialization is not an easy task; one solution is to decompose the team into multiple groups, where all agents within a group have identical behavior (1; 2; 3). Current approaches only allow the search for teams that have a static decomposition. Unfortunately, such solutions may be clearly suboptimal in domains where the contribution of each group changes with time: for example, the group of scouts is essential in a cooperative foraging scenario when no food source is known and no enemy units are detected, but for-

aging agents (respectively defending agents) are desirable once a food source is discovered or the team is under attack.

I propose to use a hierarchical architecture: a higher-level role-switch mechanism indicates which lower-level role the agent should assume in the current situation (given the agent's knowledge about the environment and its teammates). As conditions change, the agents may change roles, thus granting the team the flexibility to self-organize as appropriate to the current situation.

I employ evolutionary computation to learn such agent behaviors. Initial experiments attempt to discover role-switching mechanisms given hard-coded behaviors for roles (or vice-versa, discover behaviors for roles with hard-coded role-switching mechanisms). These experiments simplify the learning task to understand better how each of the components can be learned separately. A later set of experiments automatically learn the role-switching mechanism, as well as the behaviors for different number of roles. I also compare the proposed learning of self-organizing teams with other multiagent learning techniques as described in (4).

A last area of the thesis is concerned with the application of cooperative coevolution to learn self-organizing behaviors. I argue that the explicit decomposition of agents behaviors into roles permits a clear decomposition of learning into multiple concurrent learning processes, each concerned with a component of the entire agent behavior. This significantly decreases the number of concurrent learning processes, thus reducing the harmful co-adaptation effects and the pathologies they induce onto learning (5).

References

- J. C. Bongard. The legion system: A novel approach to evolving heterogeneity for collective problem solving. In R. Poli *et al*, editors, *Genetic Programming: Proceedings of EuroGP-2000*, volume 1802, pages 16–28, Edinburgh, 15-16 2000. Springer-Verlag.
- A. Hara and T. Nagao. Emergence of cooperative behavior using ADG; Automatically Defined Groups. In *Proceedings of the 1999 Genetic and Evolutionary Computation Conference (GECCO-99)*, pages 1038–1046, 1999.
- S. Luke. Genetic programming produced competitive soccer softbot teams for RoboCup97. In J. R. Koza *et al*, editors, *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 214–222. Morgan Kaufmann, 1998.
- L. Panait and S. Luke. Collaborative multi-agent learning: A survey. *Technical Report GMU-CS-TR-2003-01*, Department of Computer Science, George Mason University, 2003.
- R. P. Wiegand. *Analysis of Cooperative Coevolutionary Algorithms*. PhD thesis, Department of Computer Science, George Mason University, 2003.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.
Copyright 2005 ACM 1-59593-094-9/05/0007 ...\$5.00.