Layered Learning in Multiagent Systems

Peter Stone

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213-3891
pstone@cs.cmu.edu
http://www.cs.cmu.edu/~pstone

Multiagent Systems is the emerging subfield of Artificial Intelligence that aims to provide both principles for construction of complex systems involving multiple agents and mechanisms for coordination of independent agents' behaviors. As of yet, there has been little work with Multiagent Systems that require real-time control in noisy environments. Because of the inherent complexity of this type of Multiagent System, Machine Learning is an interesting and promising area to merge with Multiagent Systems. Machine learning has the potential to provide robust mechanisms that leverage upon experience to equip agents with a large spectrum of behaviors, ranging from effective individual performance in a team, to collaborative achievement of independently and jointly set high-level goals in the presence of adversaries. Learning will also help agents adapt to unforeseen behaviors on the parts of other agents, through the use of on-line adaptive methods that may include explicit opponent modelling.

My thesis will focus on learning in this particularly complex class of multiagent domains. The principal question to be answered is

> Can agents learn to work together in a real-time, noisy environment in the presence of both teammates and adversaries?

I am designing and developing a multiagent learning system in the context of robotic soccer as an example of one such domain. Based on this challenging case study, I expect to introduce a new general multiagent learning method, Layered Learning, by which similar systems can be built in any such domain. Layered Learning allows for a bottom-up definition of agent capabilities at different levels in a complete multiagent domain. Machine Learning opportunities are identified when hand-coding solutions are too complex to generate. Individual and collaborative behaviors in the presence of adversaries are organized, learned, and combined in a layered fashion. I will demonstrate the effectiveness of Layered Learning in the robotic soccer domain.

To date, I have implemented two levels of learned behaviors (Stone & Veloso 1997). First, the clients used a Neural Network to learn a low-level individual skill that allows them to control the ball effectively. Then, using this learned skill, they learned a higher-level, more "social," skill: one that involves multiple players. The second skill, the ability to estimate the likelihood that a pass to a particular teammate will succeed, was learned using a Decision Tree. For both skills, I conducted extensive empirical testing and I verified empirically that the learned skills are applicable to a game-like situation.

Although more layers are needed, the two already-implemented learned behavior levels demonstrate the feasibility of my approach. The two behaviors will allow me to continue moving upward towards high-level strategy issues, both collaborative and adversarial.

References

Stone, P., and Veloso, M. 1997. A layered approach to learning client behaviors in the robocup soccer server. *To appear in Applied Artificial Intelligence (AAI) Journal.* A shorter version, "Using Machine Learning in the Soccer Server," is available in the Proceedings of the IROS-96 Workshop on RoboCup, Osaka, Japan, November, 1996.