Layered Learning in Multi-Agent Systems

Peter Stone

December 15, 1998

CMU-CS-98-187

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3891

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Thesis Committee:
Manuela M. Veloso, Chair
Andrew W. Moore
Herbert A. Simon
Victor R. Lesser (University of Massachusetts, Amherst)

The work has been supported through the generosity of the NASA Graduate Student Research Program (GSRP). This research is also sponsored in part by the Defense Advanced Research Projects Agency (DARPA), and Rome Laboratory, Air Force Materiel Command, USAF, under agreement numbers F30602-95-1-0018, F30602-97-2-0250 and F30602-98-2-0135 and in part by the Department of the Navy, Office of Naval Research under contract number N00014-95-1-0591. Views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NASA, the Defense Advanced Research Projects Agency (DARPA), the Air Force Research Laboratory (AFRL), the Department of the Navy, Office of Naval Research, or the U.S. Government.

Keywords: Multi-agent systems, machine learning, multi-agent learning, control learning, hierarchical learning, reinforcement learning, decision tree learning, neural networks, robotic soccer, network routing

Abstract

Multi-agent systems in complex, real-time domains require agents to act effectively both autonomously and as part of a team. This dissertation addresses multi-agent systems consisting of teams of autonomous agents acting in real-time, noisy, collaborative, and adversarial environments. Because of the inherent complexity of this type of multi-agent system, this thesis investigates the use of machine learning within multi-agent systems. The dissertation makes four main contributions to the fields of Machine Learning and Multi-Agent Systems.

First, the thesis defines a team member agent architecture within which a flexible team structure is presented, allowing agents to decompose the task space into flexible roles and allowing them to smoothly switch roles while acting. Team organization is achieved by the introduction of a locker-room agreement as a collection of conventions followed by all team members. It defines agent roles, team formations, and pre-compiled multi-agent plans. In addition, the team member agent architecture includes a communication paradigm for domains with single-channel, low-bandwidth, unreliable communication. The communication paradigm facilitates team coordination while being robust to lost messages and active interference from opponents.

Second, the thesis introduces layered learning, a general-purpose machine learning paradigm for complex domains in which learning a mapping directly from agents' sensors to their actuators is intractable. Given a hierarchical task decomposition, layered learning allows for learning at each level of the hierarchy, with learning at each level directly affecting learning at the next higher level.

Third, the thesis introduces a new multi-agent reinforcement learning algorithm, namely team-partitioned, opaque-transition reinforcement learning (TPOT-RL). TPOT-RL is designed for domains in which agents cannot necessarily observe the state changes when other team members act. It exploits local, action-dependent features to aggressively generalize its input representation for learning and partitions the task among the agents, allowing them to simultaneously learn collaborative policies by observing the long-term effects of their actions.

Fourth, the thesis contributes a fully functioning multi-agent system that incorporates learning in a real-time, noisy domain with teammates and adversaries. Detailed algorithmic descriptions of the agents' behaviors as well as their source code are included in the thesis.

Empirical results validate all four contributions within the simulated robotic soccer domain. The generality of the contributions is verified by applying them to the real robotic soccer, and network routing domains. Ultimately, this dissertation demonstrates that by learning portions of their cognitive processes, selectively communicating, and coordinating their behaviors via common knowledge, a group of independent agents can work towards a common goal in a complex, real-time, noisy, collaborative, and adversarial environment.

Acknowledgements

I would like to thank many people for their support, encouragement and guidance during my years as a graduate student here at CMU.

First and foremost, this dissertation represents a great deal of time and effort not only on my part, but on the part of my advisor, Manuela Veloso. She has helped me shape my research from day one, pushed me to get through the inevitable research setbacks, and encouraged me to achieve to the best of my ability. Without Manuela, this dissertation would not have happened.

I also thank my other three committee members, Andrew Moore, Herb Simon, and Victor Lesser for valuable discussions and comments regarding my research.

Almost all research involving robots is a group effort. The members of the CMU robosoccer lab have all contributed to making my research possible. Sorin Achim, who has been with our project almost from the beginning has tirelessly experimented with different robot architectures, always managing to pull things together and create working hardware in time for competitions. Kwun Han was a partner in the software development of the CMUnited-97 team, as well as an instrumental hardware developer for CMUnited-98. Mike Bowling successfully created a new software approach for the CMUnited-98 robots. He also collaborated on an early simulator agent implementation. Patrick Riley put up with many demands and requests in a short period of time as he helped me create the CMUnited-98 simulator software. And Manuela has heroically managed all of these projects while also staying involved in the low-level details. I have been proud to share CMUnited's success with all of these people.

Many other people at CMU have helped shape my research. Members of the PRODIGY group, including Jim Blythe, Jaime Carbonell, Eugene Fink, and Xuemei Wang, encouraged me in the early states of my research. Members of Reinforcement Learning group, especially Justin Boyan, Rich Caruana, Astro Teller, Belinda Thom, and Will Uther, have also provided me with comments and suggestions along the way. Mike Bowling and Tucker Balch helped me specifically with the writing of this thesis.

Outside of CMU, I owe thanks to other members of the robotic soccer community. First, Alan Mackworth deserves special credit for originally recognizing the challenges of this domain. He and his student, Michael Sahota, were very helpful in our initial forays into this domain. Among many others, I thank Hiroaki Kitano and Minoru Asada for their endless efforts in promoting robotic soccer and making the international RoboCup competitions possible. Special thanks go to Itsuki Noda, who created the soccer server, and who has been very receptive to suggestions for its continued development. Without Noda's soccer server, my research surely would have followed a different course. I also thank Steve Chien for inviting me to work with him at JPL during the summer of 1995.

On a more personal note, I have been lucky to have many close friends who have helped to make my time as a graduate student enjoyable. Astro Teller has been a friend and confidant throughout the entire process here. Belinda Thom and Justin Boyan also stand out as people with whom I spent many enjoyable hours away from the office.

My parents, Russell and Rala Stone, and my sister Mira have always been, and continue to be, there for me at all times. And most of all, my wife Tammy has been incredibly supportive, understanding, and encouraging as she has gone through the entire graduate experience with me.

Contents

List of Figures

List of Tables