Research Summary and Plans

Peter Stone

The field of Artificial Intelligence is at a crossroads. Much of the past research in the field has focused either on high-level reasoning from abstract, ungrounded representations; or on interpreting raw sensor data towards building grounded representations. However, neither of these foci taken alone is sufficient for deploying, practical, real-world AI systems. In recent years, an active contingent within the field has focused on creating *complete* autonomous agents: those that sense their environment; engage in high-level cognitive decision-making; and then execute their actions in the environment.

As the field progresses in this direction, individual autonomous agents, either in software or physically embodied, are becoming more and more capable and prevalent. Multiagent systems consisting of homogeneous or similar agents are also becoming better understood. However, to successfully interact in the real world, agents must be able to reason about their interactions with heterogeneous agents of widely varying properties and capabilities. My ongoing and past research over the course of a decade has played a key role by focusing on particularly challenging multiagent problems involving heterogeneous agents and pushing the boundary of possibility by applying and developing cutting-edge machine learning techniques. My main approach is to study complete agents in specific, complex environments, with the ultimate goal of drawing general lessons from the specific implementations.

The first domain in which my research has focused is *robot soccer*. Robot soccer pits two teams of independently-controlled agents against each other. In 1994, as a Ph.D. candidate, I founded the robosoccer project at Carnegie Mellon University, along with my advisor, Manuela Veloso. I was also a founding member of the RoboCup executive committee, and am a member of the board of trustees of RoboCup, an organization that now attracts more than 1000 researchers to its annual competitions and conferences. As a participant, I have helped create teams of agents that have won 5 RoboCup championships in 3 different RoboCup leagues. The technical contributions related to my research in the robot soccer domain have included a new multiagent machine learning algorithm; a flexible team structure for domains in which teammates can periodically synchronize; a hierarchical machine learning paradigm in which learning at lower layers directly facilitates learning at higher layers; and a communication protocol for single-channel, low-bandwidth communication environment; among others. With the support of a recent NSF CAREER award, my ongoing research in this domain includes the development of a legged team using Sony Aibo robots, as well as a coaching approach that allows an omniscient coach agent to *learn* the appropriate strategic advice to give to its team members based on the observed behaviors of the current opponent. I am also using the robot soccer domain as a tool for motivating students, both undergraduates and graduates, to become immersed in AI research challenges by actively incorporating it into my classes.

The second main application domain of my past research is *autonomous bidding agents*. Motivated by the trading agent competition hosted by the University of Michigan, in which an agent acting as a travel agent buys flights, hotel rooms, and entertainment tickets for its customers, I created an adaptive agent that changes its bidding strategies based on the results of previous auctions. This domain is particularly challenging due to the fact that the value of each good is closely tied to the evolving prices of all the other goods (complementarity and substitutability), which are being sold simultaneously. Our agent finished in first place in the 2000, 2001, and 2003 trading agent competitions. The technical contributions behind this agent include a general decision-theoretic approach to bidding in simultaneous auctions; and a new boosting-based machine learning algorithm for conditional density estimation. I used this new algorithm to successfully predict the closing prices of future auctions. As a practical offshoot of this research, while a senior researcher at AT&T Labs — Research, I then created agents to help inform AT&T's bidding in the FCC spectrum auction of December 2000, an auction that brought in over $16 billion dollars. We discovered a novel, successful bidding strategy in this domain that allows the bidders to increase their profits significantly over a reasonable default strategy. My ongoing research pertaining to autonomous bidding agents focuses mainly on a supply chain scenario in which agents learn how best to procure components, manage a manufacturing process, and respond to customer demand in a computer-manufacturing scenario.

While continuing my research in these domains, my current research goals include transitioning my successful research results to new challenging multiagent domains. In particular, with the aid of funding from NASA, I am working on applying one of my main thesis contributions, layered learning, to the control of humanoid robots for space applications; and with the aid of an IBM faculty award, I am making progress towards using layered learning to bring IBM's autonomic computing initiative to fruition. For entire systems to be able to self-diagnose failures and repair themselves, there will need to be learning components at multiple levels, including the OS, databases, and networking modules. Thus a paradigm like layered learning will be essential to keeping the entire system operating smoothly.

In summary, the most exciting research topics to me are those inspired by challenging real-world problems. Furthermore, successful research results involve fully implemented solutions. I have addressed AI challenges including deliberative planning, real-time multiagent collaboration, machine learning, and e-commerce, producing concrete systems both in simulation and on robots. My ongoing research aims to bring these contributions to still more practical, real-world problems.