

Introduction to the Special Issue

Ten Years of Autonomous Agents and Multiagent Systems

Liz Sonenberg, Peter Stone, Kagan Tumer, Pinar Yolum

■ *The articles in this special issue have been specifically commissioned to provide a snapshot of current activity in the autonomous agents and multiagent systems communities.*

The research addressed in the autonomous agents field covers a wide spectrum of levels from the cognitive to the organizational, exploits diverse mechanisms and approaches, and has had a major impact on many aspects of artificial intelligence research. In 2011 the Autonomous Agents and Multiagent Systems (AAMAS) conference series celebrated its 10th anniversary, having begun as the successful merger of three related events that had run for some years previously. AAMAS 2002, the first of the new events, was held in Bologna, followed by Melbourne (2003), New York (2004), Utrecht (2005), Hakodate (2006), Honolulu (2007), Estoril (2008), Budapest (2009), Toronto (2010), and Taipei (2011).

The 2011 AAMAS conference received 575 submissions, and 126 papers were selected for publication as full papers. Representation under all submissions of topics (measured by first keyword) was broad, with top counts in areas such as teamwork, coalition formation, and coordination (31), distributed problem solving (30), game theory (30), planning (26), multiagent learning (24), and trust, reliability, and reputation (17). The tag cloud (figure 1), generated from the titles of the full papers at the conference, conveys a sense of the relative prominence of topics.

The articles in this special issue have been specifically commissioned to provide a snapshot of current activity in the



Figure 1. The Tag Cloud Conveys a Sense of the Relative Prominence of Topics.

autonomous agents and multiagent systems communities. Drawing on our experiences as general and program cochairs of the 2011 AAMAS conference, our request to the authors selected for this special issue was to include material motivating the topic; to give an overview of research in the subfield with pointers to contemporary review articles where feasible; to discuss some of the authors' own contributions to the topic; and to offer an account of challenges or future directions. Of course, none of these elements could be covered comprehensively, so the reader should view this collection as providing an accessible perspective on the diversity of the field, presented through the voices of leading researchers. We hope that readers unfamiliar with the field will be enticed to explore further, and readers closer to the agent research communities will appreciate these articles as providing one of many possible snapshots of a highly active and productive research field.

A feature of research in agent systems is integration. Research on individual agents typically explores the integration of multiple components and capabilities, such as planning, learning, reactivity, goal orientedness, to yield action, typically while the agent is situated in a complex environment. Multiagent systems are distributed systems of actors—robots, software, or humans—that interact, in a cooperative or competitive way, and achieve outcomes through the explicit or implicit integration of their capabilities. Accordingly, each of the following papers offers a perspective on integration—though the foci and the underlying mechanisms are very varied.

The opening article in this collection, An Overview of Recent Application Trends at the AAMAS Conference: Security, Sustainability, and Safety by Manish Jain, Bo An, and Milind Tambe,

profiles work that ties agent systems research into practical applications. The authors describe several key research threads that draw on the framework of game theory and present deployed applications that are use inspired and advance the state of the art. The article provides an account of several security applications, including relevant algorithms, evaluation, and challenges, and touches on applications in energy management, disaster management, and traffic management.

The second article, Agent-Based Modelling and Simulation by Franziska Klügl and Ana Bazzan, also uses traffic management as an application domain, but in this case, to illustrate agent-based modeling and simulation (ABMS) techniques. This article introduces the core idea of ABM as being to model active components in a system as agents, interacting with some of their peers as well as with the environment. Running simulations then produces empirical data that can be analyzed relative to natural phenomena. The article discusses ABMS as compared with other modeling and simulation approaches, mentions several toolsets, provides pointers to further work in several domains, and concludes by describing several challenges.

Challenges are the central theme in the article, Multiagent Learning: Basics, Challenges, and Prospects by Karl Tuyls and Gerhard Weiss. The infeasibility of engineering agents so that they can anticipate all appropriate behaviours, provides the rationale for designing agents able to adapt their behaviour from experience— from the environment and from interaction with other agents. This paper offers a bird's-eye perspective on multiagent learning, that is, the field that integrates machine-learning techniques in multiagent system, by reflecting on the nature, foundations, and evolution of the field. Pointers are provided to recent

survey articles, a brief introduction is offered of key algorithms, and the article concludes with an identification of challenges in multiagent learning and promising future directions.

Unsurprisingly, many of the intellectual roots of multiagent systems research can be found in the field of distributed problem solving, and that subfield continues to be very active. The article Distributed Problem Solving by William Yeoh and Makoto Yokoo focuses on distributed constraint reasoning topics such as distributed constraint satisfaction problems (DCSPs) and distributed constraint optimization problems (DCOPs). Supported by a running example involving a wireless LAN scenario, the authors discuss major classes of CSP, DCSP, COP, and DCOP algorithms. They also provide links to a set of sample testbeds and an introduction to recent extensions to distributed constraint reasoning models that seek more faithfully to capture real-world problems.

A call to arms to the research community to invest in research using physical robots forms the passionate core of the article I Have a Robot and I'm Not Afraid to Use It! Gal Kaminka reminds us that a recurring theme of agent research is the focus on integrating capabilities, including learning and strategic decision making, in rich environments, and draws parallels with the systems thinking demands of the robotics research communities. By touching on individual architectures, and then offering an account of successful cross-fertilization between research on teamwork in multiagent systems and on multirobot systems, this article highlights both success stories to date, and opportunities for the future.

Negotiating Agents by Catholijn Jonker, Koen Hindriks, Pascal Wiggers, and Joost Broekens profiles another active subfield of the AAMAS community. The article makes mention of the design of fully automated negotiating agents, but places more emphasis on aspects required to build systems to support human negotiators. Drawing on experiences from the Automated Negotiating Agents Competition (ANAC) held in conjunction with AAMAS in 2010 and 2011, the authors describe state-of-the-art negotiation support approaches and key challenges in this field, a field where the ambition is to create technologies to allow hybrid human-agent negotiating teams to outperform human negotiators.

Focusing on theory, the final article brings us back to one of the early and highly active subfields of agent systems research. In Logics for Multiagent Systems, Wiebe van der Hoek and Michael Wooldridge offer a narrative that motivates logic as a tool for reasoning about multiagent systems, touches briefly on the basic logics used, and then discusses the role that logics might play in the development of agents. The authors structure the

presentation around two current strands of research — cognitive models of rational action, and models of the strategic structure of systems. Accordingly, the reader will meet logics that support individual agent specification — intention logic and BDI logic — and will also see logics of cooperative ability — coalition logic and alternating-time temporal logic. The article contains both an accessible narrative for newcomers to the area and in depth technical content for more advanced readers.

In aggregate, these invited papers, authored by a selection of highly credentialed authors, provide a sense of the liveliness of a highly influential, and still growing, subfield of artificial intelligence. We hope you enjoy reading these papers, and that you find yourself coming back to them more than once, to gain further insights, and for pointers to further reading.

Liz Sonenberg (l.sonenberg@unimelb.edu.au) is a professor of information systems in the Department of Computing and Information Systems and pro vice-chancellor (research collaboration) for the University of Melbourne, Australia. Her research targets the design of reasoning machinery for systems that exhibit complex collaborative behaviors with a focus on teamwork and on human-agent interaction. Sonenberg was general cochair of the 2011 AAMAS Conference.

Peter Stone (pstone@cs.texas.edu) is an Alfred P. Sloan research fellow, a Guggenheim fellow, and a professor in the Department of Computer Sciences at the University of Texas at Austin, USA. He is the founder and director of the Learning Agents Research Group (LARG) within the Artificial Intelligence Laboratory and his research focuses mainly on machine learning, multiagent systems, and robotics. Stone was the general cochair of the 2011 AAMAS Conference.

Kagan Tumer (kagan.tumer@oregonstate.edu) is a professor of robotics and control, in the School of Mechanical, Industrial, and Manufacturing Engineering at Oregon State University, USA. His main research interests are in designing control and optimization algorithms for large complex systems, with particular interests in learning and coordination. Tumer was a program cochair of the 2011 AAMAS Conference.

Pinar Yolum (pinar.yolum@boun.edu.tr) is an associate professor in the Department of Computer Engineering, Bogazici University, Turkey. Her current research interests in multiagent systems include trust, negotiation, and commitments, with broader activities in service-oriented computing and the design and analysis of protocols. Yolum was a program cochair of the 2011 AAMAS Conference.