

# The Workshop Program at the Nineteenth National Conference on Artificial Intelligence

*Ion Muslea, Virginia Dignum, Daniel Corkill, Catholijn Jonker, Frank Dignum, Silvia Coradeschi, Alessandro Saffiotti, Dan Fu, Jeff Orkin, William Cheetham, Kai Goebel, Piero Bonissone, Leen-Kiat Soh, Randolph M. Jones, Robert E. Wray, Matthias Scheutz, Daniela Pucci de Farias, Shie Mannor, Georgios Theodorou, Doina Precup, Bamshad Mobasher, Sarabjot Singh Anand, Bettina Berendt, Andreas Hotho, Hans Guesgen, Michael T. Rosenstein, and Mohammad Ghavamzadeh*

■ AAAI presented the AAAI-04 workshop program on July 25–26, 2004 in San Jose, California. This program included twelve workshops covering a wide range of topics in artificial intelligence. The titles of the workshops were as follows: (1) Adaptive Text Extraction and Mining; (2) Agent Organizations: Theory and Practice; (3) Anchoring Symbols to Sensor Data; (4) Challenges in Game AI; (5) Fielding Applications of Artificial Intelligence; (6) Forming and Maintaining Coalitions in Adaptive Multiagent Systems; (7) Intelligent Agent Architectures: Combining the Strengths of Software Engineering and Cognitive Systems; (8) Learning and Planning in Markov Processes—Advances and Challenges; (9) Semantic Web Personalization; (10) Sensor Networks; (11) Spatial and Temporal Reasoning; and (12) Supervisory Control of Learning and Adaptive Systems.

AAAI presented the AAAI-04 workshop program on Sunday July 25 and Monday, July 26, 2004 at the San Jose McEnery Convention Center and the adjacent headquarters hotel in San Jose, California. This program included sixteen workshops covering a wide range of topics in artificial intelligence. Most workshops were held on a single day, and atten-

dance was limited to approximately 25 to 65 participants. Participation at these workshops was by invitation from the workshop organizers. The titles of the workshops were as follows: (1) Adaptive Text Extraction and Mining; (2) Agent Organizations: Theory and Practice; (3) Anchoring Symbols to Sensor Data; (4) Challenges in Game AI; (5) Fielding Applications of Artificial Intelligence; (6) Forming and Maintaining Coalitions in Adaptive Multiagent Systems; (7) Intelligent Agent Architectures: Combining the Strengths of Software Engineering and Cognitive Systems; (8) Learning and Planning in Markov Processes—Advances and Challenges; (9) Semantic Web Personalization; (10) Sensor Networks; (11) Spatial and Temporal Reasoning; and (12) Supervisory Control of Learning and Adaptive Systems.

## Adaptive Text Extraction and Mining

The terabytes of available online texts contain invaluable information for virtually every domain of activity. However, existing tools for accessing and exploiting this data fall short of users'

needs. Recently, researchers have become quite interested in this problem and they have made significant progress in developing techniques for the automatic extraction and mining of text corpora. Adaptive text extraction and mining (ATEM) is an extremely active area of research that lies at the intersection of diverse fields such as information extraction, text mining, machine learning, data mining, natural language processing, and information integration. The purpose of the ATEM-2004 workshop—the fifth such meeting since 1999—was to bring together researchers and practitioners from these communities to discuss recent results, trends, and open problems. More than thirty people attended the workshop from both academe and industry. It featured an invited talk, seven full-length papers, nine posters, and a panel. The workshop's invited presentation by Mark Craven (University of Wisconsin) focused on the uses of information extraction in biomedical domains. Craven discussed at length two primary challenges for the ATEM community: annotating genomes and proteomes based on the available online literature, and annotating high-throughput experiments (such as providing an explanation of why a set of genes or proteins exhibits a particular behavior). A common ATEM theme was the learning (supervised or unsupervised) of extraction rules for a particular corpus or domain. However, a number of papers explored a variety of other ideas. For example, the workshop featured three papers on KnowItAll, a domain-independent system developed at the University of Washington that extracts massive amounts of information from the web in an automated, open-ended manner. Two other papers focused on the use of ontologies in learning for information extraction. Some of the participants described techniques related to or borrowing from natural language processing, such as web site summarization, shallow parsing, and lexical semantics domain models. The workshop ended with a lively discussion that began with a critical survey on the evaluation of information extraction systems, followed by a panel on open problems and future trends in ATEM. Among the topics debated by the panel were the need for

better unsupervised information extraction techniques, the use of information extraction in search engines, the advantages and disadvantages of using a standard repository of extraction tasks, and embedding ATEM tools in real world application. The papers from this workshop were published as AAAI Technical Report WS-04-01, and are available from AAAI Press.

Ion Muslea, *SRI International*

## Agent Organizations: Theory and Practice

The notion of an agent organization is becoming a very important in multiagent systems (MAS) research. The notion is based on the fact that many multiagent systems are open (agents can enter and leave) and contain many heterogeneous agents that must coordinate their efforts. The notion of "agent organizations" means that the organizational and individual (agent) perspective has to be integrated. We need explicit representations and implementations of the organizational aspects of the MAS. As systems grow to include hundreds of agents we must move from an agent-centric view of coordination and control to an organizational-centric view. Moreover, these organizations must be able to adapt themselves to a changing environment. Based on previous workshops on social notions of multiagency and on the use of MAS approaches to the modeling and simulation of organizations and societies, this workshop brought together researchers from the organizational modeling and MAS research fields. The primary themes of the workshop were organizational models, social issues in organizations, and adaptation. Within the first theme, we focused our attention on ways of describing organizations for MAS and how the organizational structure influences the performance of the agents. Social issues that were discussed under the second theme included ways that cliques are formed (and stay together), what are determining factors, and how simulation models can help to understand the emergence of cliques in groups of people. Another issue that arose was the models that agents should have of social phenomena to af-

fect joined behavior. Ways for agents and organizations to adapt themselves to a changing environment were discussed. The workshop was a great success, as it showed the relation between several issues on organizational modeling and paved the way for more structural collaboration on the central themes of the workshop; it will therefore certainly be continued in some form in the future. The papers from this workshop were published as AAAI Technical Report WS-04-02, and are available from AAAI Press.

Virginia Dignum, *University of Utrecht*;  
Daniel Corkill, *University of Massachusetts*;  
Catholijn Jonker, *Free University Amsterdam*; and Frank Dignum, *University of Utrecht*

## Anchoring Symbols to Sensor Data

*Anchoring* is the problem of how to create, and maintain in time, the connection between symbol- and signal-level representations of the same physical objects. Anchoring is an important aspect of the connection between symbolic reasoning and sensory processes in autonomous robots. Anchoring must also occur in multiple robot systems whenever the robots exchange information via symbolic representations. It is also needed for efficient human-robot cooperation. This workshop was the third in a series of meetings and publications devoted to perceptual anchoring. During the workshop two invited talks and ten technical papers were presented. The technical contributions indicated the growing maturity of the anchoring community, which is moving away from simple cases and beginning to tackle more complex issues. Examples of complex issues include dealing with indistinguishable objects (John F. Santore and Stuart C. Shapiro, University of Buffalo) and failures in the anchoring process (Mathias Broxvall, Silvia Coradeschi, Lars Karlsson, and Alessandro Saffiotti, Örebro University). Several participants proposed the use of learning and a bottom-up approach to anchoring, including Joseph Modayil and Benjamin Kuipers (University of Texas), Derek Magee, Chris J. Needham, Paulo Santos, Anthony G. Cohn, and David C. Hogg

(University of Leeds), and Eric Hoernagl, Patrick Poelz, and Erich Prem (Austrian Research Institute for Artificial Intelligence).

The search for a suitable general framework for anchoring continued at this workshop. That search has moved toward the use of more general structures, inspired by work in cognitive science (Antonio Chella, University of Palermo; Silvia Coradeschi, Örebro University; Marcello Frixione, University of Salerno; and Alessandro Saffiotti, Örebro University). It is also focusing on the computational, representational and algorithmic aspects of the framework (Fredrik Heintz and Patrick Doherty, Linköping University). Some more philosophically oriented contributions were also presented at the workshop. They addressed the problem of understanding the nature of a "symbol" (Erich Prem), taking insights from neurophysiology (Angel del Pobil, Enric Cervera, and Eris Chinellato, Universitat Jaume I), and discussing frameworks to model and understand the brain (Andreas Knoblauch, Rebecca Fay, Ulrich Kaufmann, Heiner Markert, and Günther Palm, University of Ulm). One of the goals of this workshop was to make contact with other related communities that address tasks in which anchoring is present even if it is not explicitly mentioned. The two invited speakers were chosen with this goal in mind. The first speaker (Jeffrey Siskind) came from the vision community; the second speaker (Gerhard Sagerer) is at the intersection between natural language, vision, and robotics. Both speakers clearly outlined the need to address the anchoring problem in these communities. Lively discussions characterized the entire workshop, indicating that there is increasing excitement about clarifying and identifying the difficult issues of the anchoring problem, both from a practical and a foundational point of view. Discussions also focused on finding viable solutions to those issues for the purpose of building autonomous systems. The papers from this workshop were published as AAAI Technical Report WS-04-03, and are available from AAAI Press.

Silvia Coradeschi and Alessandro Saffiotti, *University of Örebro*

## Challenges in Game Artificial Intelligence

More than sixty people attended the two-day Challenges in Game AI workshop, which aimed to bridge the gap between the game industry and academe, and brought together game AI experts from seven game-development studios and fourteen universities, plus others from the military and middleware.

Day one began with a presentation about the lessons learned in building a generic AI engine, from the perspective of BioGraphic Technologies' chief technology officer, Paul Kruszewski. Duane Szafron demonstrated and spoke about the University of Alberta's ScriptEase code generation tool for Neverwinter Nights. Hector Munoz-Avila discussed applying hierarchical task network planning to Unreal bots at Lehigh University. Researchers from the University of Southern California's Institute for Creative Technologies presented an interface for tactical and strategic AI used in Full Spectrum Command. Radical Entertainment's Ben Geisler used his thesis research of integrating machine learning into *Soldier of Fortune II* as a starting point for an interactive discussion about appropriate AI technologies for games. The day wrapped up with an entertaining commentary on bad game AI by Baylor Wetzel.

Day two kicked off with case studies of three commercial games: *Mythica*, *The Suffering*, and *F.E.A.R.* Carnegie Mellon graduate student Curt Bererton demonstrated the use of particle filters for agent navigation in the *Crystal Space* engine. Researchers from the U.S. Army MOVES Institute discussed techniques for finding dynamic cover positions in *America's Army*. Radical Entertainment's Borut Pfeifer talked about creating emergent game play, in terms of AI and game design. Finnegan Southey presented game-play analysis tools in development at the University of Alberta, which have been utilized by Electronic Arts in *FIFA 2005*. The workshop concluded with Michael Mateas's (Georgia Institute of Technology) talk about his work on *Façade*, followed by a panel discussion about interactive narratives.

Aside from the presentations, participants commented that highlights were

getting to meet a diverse group of talent from the game industry and academia, and honest discussions about strong AI versus entertainment AI. The papers from the workshop were published as AAAI Technical Report WS-04-04, and are available from AAAI Press.

Dan Fu, Stottler Henke, and Jeff Orkin,  
*Monolith Productions*

## Fielding Applications of Artificial Intelligence

"Knowledge unapplied in action is a pale academic vanity" (Francis Bacon).

The Fielding Applications of AI workshop concluded that the process of fielding a real-world AI application represents about seventy percent of the total time needed from conception to use in a production environment. Since the benefits of AI systems can only be realized when they are actually used to solve problems, the importance of transitioning systems from conception to production cannot be understated. The talks at the workshop highlighted many of the benefits and best practices of fielding AI applications.

Helen Hastie (Lockheed Martin Advanced Technology Laboratories) pointed out that establishing users' trust in a fielded system is a key issue in the success of a system. She presented a spoken-language understanding system for military environments called "SUSIE." The military users of the system would be provided with answers to their spoken enquiries. This meant creating capabilities like remembering past information requests in order to speed communication (that is, handling questions like "SUSIE, what about now?" when a specific question was asked previously).

Bogdan Stanescu and Marcel Barbulescu (George Mason University) discussed best practices that impact the success of a fielded AI system. These factors include having a good business case, working closely with users during the entire project development, making the system easy to use and maintain, following formal development and deployment practices, immediately supporting any user questions or needs, and managing user expectations.

Debra Schreckenghost (NASA Johnson Space Center) noted that humans'

ability to efficiently and conveniently interact with AI systems would become more important to AI systems. For example, the system may need to track completion status of tasks performed by a human without having the human provide this information.

Other presentations showed additional benefits and best practices, such as (1) fielding is a good way to evaluate an AI technique; (2) in attempting to solve a real-world problem other areas to apply AI can be discovered; (3) researchers should apply methodologies like software engineering and Six Sigma; and (4) there can be a financial benefit to the researcher from fielding applications.

Progress in the field of AI takes place in many ways, including theoretical and fielded research. However, given the time and skills needed to field a system, many AI researchers choose not to do so. While the process of learning during the creation is certainly useful to the researchers themselves, there may still be an overall acceptance problem for the field of AI because of the lack of demonstrable quantifiable benefit. It was suggested that if a larger percentage of researchers transition their work to a real-world application, quicker progress in generating scientific results could be obtained. Consequently, the AI community should find ways to encourage and support researchers to field their work. This could in part be facilitated by sharing knowledge and experience about the skills needed to successfully transition theoretical work to solving real problems. This workshop saw itself as one step in that direction. The papers from the workshop were published as AAAI Technical Report WS-04-05, and are available from AAAI Press.

William Cheetham, Kai Goebel, and Piero Bonissone, *General Electric Global Research*

## Forming and Maintaining Coalitions and Teams in Adaptive Multiagent Systems

This workshop focuses on the issues of coalitions in dynamic multiagent systems: specifically, on issues surround-

ing the formation of coalitions among possibly self-interested individuals, and on how coalitions adapt to change in dynamic settings through the choices of individual members. Traditionally, an agent with complete information can rationalize to form optimal coalitions with its neighbors for the purpose of problem solving. However, in a noisy and dynamic environment, information cannot be relayed among the agents frequently enough, centralized updates and polling are expensive, and the supporting infrastructure may partially fail. Agents are thus forced to form suboptimal coalitions. Similarly, changes in the environmental dynamics may invalidate some of the reasons for the original existence of a coalition. In this case, individual agents may influence the objectives of the coalition, encourage new members and reject others, and the coalition as a whole adapts as a larger organism. Therefore, agents need to reason to form a successful coalition rather than an optimal one, and to maintain the coalition to suit its changing needs. The workshop had a collection of eight papers, covering issues on team and coalition formation with a broad range of approaches for its machine learning and team or coalition evaluation strategies.

John Anderson, Brian Tanner, and Jacky Baltes (University of Manitoba) presented a reputation profile based on a memory model in the robotic-soccer domain, in which agents learn to team up with better agents, and agents learn to improve themselves to gain better reputation.

Dipyaman Banerjee, Sabyaschi Saha, and Sandip Sen (University of Tulsa) described an expertise-selection mechanism that adapts to market needs, in which an agent predicts for itself the best expertise to possess and changes accordingly, with an adaptive exploration strategy for revising its expertise.

Xin Li and Leen-Kiat Soh (University of Nebraska-Lincoln) discussed case-based reinforcement learning for coalition planning and instantiation, in which agents learn how to plan for a coalition and how to instantiate it through utility-tagged cases in a noisy and uncertain environment.

The paper by Priyang Rathod and Marie desJardins (University of Mary-

land) presented an employer-employee-company model for stable team formation, in which an agent can form a company, hiring other agents to form the best team possible and bid for a task, decide to keep the company intact based on the utility, and make appropriate offers. The agent can also move from one company to another based on the offers it receives (for example, shares, commission, dividend, or penalty).

Leen-Kiat Soh (University of Nebraska-Lincoln) presented a cooperative learning team model adopted from research in education, in which agents work behind the scenes to help teachers teach better through automated profiling of students and help students learn better through automated "buddy group" formation based on the assignments at hand and the observed student-characteristics profiling.

Finally, the presentation by R. Sorbello and A. Chella (Università di Palermo) and R. C. Arkin (Georgia Institute of Technology) discussed a coalition formation framework based on a political system for voting and role assignment, applied to mission-based robots to search for bombs, in which the direction of the coalition is determined through forming the dominant coalitions among political parties (of agents) of different tendencies.

The papers from the workshop were published as AAI Technical Report WS-04-06, and are available from AAI Press.

Leen-Kiat Soh, *University of Nebraska-Lincoln*

## Intelligent Agent Architectures: Combining the Strengths of Software Engineering and Cognitive Systems

The AAI 2004 workshop on Intelligent Agent Architectures focused on issues in the development of intelligent agent architectures that combine the functional structures and mechanisms of cognitive architectures with software engineering principles that have been applied in various types of agent systems and traditional software systems.

This workshop attracted a variety of participants with backgrounds in intelligent agents, cognitive architectures, and (to a lesser extent) software engineering. Contributed presentations covered topics including component-oriented analyses of architectures, architectural integrations, applications, and comparative frameworks and methodologies.

Four invited presentations framed the workshop discussions. Pat Langley (Center for the Study of Language and Information) discussed the important tensions between engineering concerns and the necessary integration of interdependent components into a cognitively oriented architecture. Manuela Veloso (Carnegie Mellon University) described her development experiences with a number of large and successful robotic behavior systems. John Laird (University of Michigan) reiterated the importance of implementation-level interdependencies in agent architectures, using Soar as a case study. Finally, Ian Horswill (Northwestern University) discussed the advantages and disadvantages between rigid hierarchical views of architectures and more fluid views based on interacting components.

One important issue that arose in the discussion was the need to find useful levels of abstraction and presentation for agent architectures, and, in particular, to distinguish between agent architectures and agent languages. Architecture design must carefully weigh features versus performance because new features can degrade overall performance. In contrast, language design encourages expansion of the feature set, making it easier to represent and communicate solutions. Thus, in applying lessons from software engineering, we must take care to distinguish between performance enhancements and usability enhancements. Architectural evolution should be driven by performance improvement.

We hope future gatherings of researchers in these areas will continue to focus on more mature incarnations of frameworks and implementations of both high-level agent languages and their corresponding architectural substrates. High-level frameworks will

hopefully provide a common language for future exploration, implementation, and comparison of well-engineered agent systems.

The papers from the workshop were published as AAAI Technical Report WS-04-07, and are available from AAAI Press.

Randolph M. Jones, *Colby College*;  
Robert E. Wray, *Soar Technology*; and  
Matthias Scheutz, *University of Notre Dame*

## Learning and Planning in Markov Processes — Advances and Challenges

The Learning and Planning in Markov Processes workshop brought together a large and diverse population of researchers whose research focuses on Markov decision processes (MDPs). The nineteen accepted papers and the three invited talks covered a spectrum of topics and included both theoretical advances and applications to real-world problems.

In the first of the three invited talks, entitled “Learning and Pricing,” Martin Puterman from the University of British Columbia addressed the problem of how to change prices over time to maximize expected revenues in the presence of unknown demand distribution parameters. In the second talk, “Knowledge Representation in TD Networks,” Richard S. Sutton from the University of Alberta talked about temporal difference (TD) methods and how they bring us closer to the goal of representing world knowledge in entirely predictive terms, grounded in observations and actions. In the third talk, “Optimizations in Large Stochastic Transportation Networks,” Milos Hauskrecht from the University of Pittsburgh discussed and proposed solutions to a couple of optimization problems arising in large stochastic optimization networks. In particular, he addressed the problem of controlling transportation flows and allocating resources with optimal coverage over an unreliable network topology.

Several of the accepted papers were presented in two oral presentation tracks, one more theoretical and one more applied. There was also a final

poster session, giving participants an opportunity to learn about and discuss all accepted work. Workshop slides, supplementary notes, and continuing discussions can be accessed on the workshop web site. The papers from the workshop were published as AAAI Technical Report WS-04-08, and are available from AAAI Press.

Daniela Pucci de Farias, and Georgios Theodorou, *Massachusetts Institute of Technology*; Shie Mannor, *McGill University*; and Doina Precup, *McGill University*

## Semantic Web Personalization

The primary goal of the Semantic Web Personalization workshop was to focus on the effective integration of semantic knowledge with standard web mining and machine-learning techniques in order to create more intelligent web personalization and recommender systems. The notion of “semantic knowledge” was interpreted broadly to include both the extraction of semantic structures from one or more web sites, as well as the use of existing or derived ontologies in the context of semantic web. Current recommender systems, while useful, cannot recommend items or provide personalized content based on the complex semantic relationships and properties of objects in the underlying domain. The effective use of this type of semantic knowledge is essential in enabling systems to automatically explain or reason about the user models or recommendations.

The workshop was successful in bringing together researchers and practitioners from different communities—in particular, semantic web and web intelligence—who are developing methods and systems that use semantic structures in the web and other AI techniques for personalization. The workshop call for papers generated contributions from eight countries in North and South America, and from Europe. The nine accepted papers covered a wide range of issues falling into three broad categories: architectures and systems; data and knowledge modeling, integration, and management; and enabling technologies.

Overall, the contributions and the

discussion showed that personalization systems profit from standardized content ontologies (such as ACM topic classification) and from schemes developed within and by user communities themselves. The importance of the distributed, heterogeneous, and decentralized nature of web content was stressed, and solutions were presented for interlinking existing ontologies and for learning new relationships about dynamically evolving communities. Other forms of semantic annotations of content, and of learning from usage (for example, viewed pages and viewing time) and explicit user input (for example ratings) were also used for personalization and recommendation. Highlighted applications included e-commerce recommender systems, research literature annotation and management, and learning platforms. Several contributions proposed novel recommendation techniques, which may be suitable for seamless integration of knowledge from multiple sources, including content, structure and usage.

The workshop generated a significant amount of interest and discussion, which identified a number of relevant areas for future research. Concerning the question of how personalization should be done, it was pointed out that it is crucial to explain recommendations, that is, to provide additional information for supporting a user’s final decision-making. This is a key factor for creating trust in recommender agents. Another relevant property, which is expected to profit from integrating semantics, is serendipity—the generation of nonobvious recommendations. Personalization can be performed on the semantic web (and may prove to be one of its most successful applications), and it can help build the semantic web, for example, when users provide personalized ontologies. Furthermore, evaluation issues remain central for personalization. This includes both metrics as well as experimental methodology. Three questions were proposed to assess the extent to which a user’s knowledge gap has been closed: How likely is it that the user would have accessed the recommended item anyway? What does she do with the recommendation? And how satisfied is she with the rec-

ommendations? These three questions can lead to an evaluation process that combines and develops accuracy-based measures and measures of usefulness.

The papers from the workshop were published as AAAI Technical Report WS-04-09, and are available from AAAI Press.

Bamshad Mobasher, *DePaul University*; Sarabjot Singh Anand, *University of Ulster*; Bettina Berendt, *Humboldt University*; and Andreas Hotho, *University of Kassel*

## Spatial and Temporal Reasoning

The AAAI-2004 workshop on Spatial and Temporal Reasoning continued a series of workshops that started eleven years ago at the 1993 International Joint Conference on Artificial Intelligence in Chambéry, France. The common goal of these workshops was to bring together related communities of researchers that study representing and reasoning about either space or time—or both. However, this year for the first time in the history of the workshops, one of the cochairs was missing. Frank Anger, who inspired and launched many of these workshops, died in a tragic car accident on July 7. His wife, Rita Rodriguez, who was edited the workshop proceedings, was critically injured in that accident.

Compared with previous workshops, the AAAI-2004 workshop on Spatial and Temporal Reasoning was a small one. Frank and Rita's absence from the workshop was certainly a reason for that, as in the past they were always able to attract additional attendees. Another reason was that a similar workshop was held at the European Conference on Artificial Intelligence (ECAI-2004) a few weeks later, which drew away some of the potential attendees. Nevertheless, the participants in the workshop had an inspiring time with interesting presentations and fruitful discussions.

Unlike previous workshops, the papers in 2004 focused primarily on temporal issues, rather than spatial ones. They fell into two categories. The first category of papers dealt with models of time and space, and reasoning about these models, whereas the second cat-

egory of presentations discussed constraint-based temporal reasoning techniques. The authors of the presentations came from around the world, including the USA, France, Germany, and New Zealand. Although each paper was of interest on its own, the two discussion sessions accompanying the presentations of the papers were of even greater value.

Hans Guesgen, *University of Auckland*

## Supervisory Control of Learning and Adaptive Systems

With human supervisory control, a human operator intermittently takes control of a process that is otherwise controlled by a computer. Supervisory control involves both autonomy and intelligence, although the latter is normally attributed solely to the human operator. One goal of this workshop, therefore, was to bring together researchers in various areas to explore the role of supervisory control for AI systems. We were fortunate to have high-quality presentations on human-computer interaction, multimodal interfaces, robotics, teleoperation, natural language, machine learning, and assistive technology. Common ground was easy to find, as evidenced by the lively discussions throughout the day.

During his keynote talk, Paul Schenker from the Jet Propulsion Laboratory described many opportunities for AI research to influence future NASA missions. In broad terms, autonomous systems will play key roles for in-space assembly, maintenance, and servicing operations, and for habitation and exploration of planetary and lunar surfaces. But unlike the common "all-or-none" notion of autonomy, autonomy for space systems implies supervisory control—with a human in the loop to improve safety and overall performance of the automated system.

The invited talk by Ken Goldberg (University of Berkeley), with Dana Plautz (Intel Research), described the evolution of several systems for web-based human-robot interaction. For supervisory control of robots, web interfaces not only have unique advantages

and challenges but also present interesting possibilities to study multiuser interaction. For instance, one can study a group's dynamics and, in particular, how leaders and followers emerge as individuals cooperate and compete for control of a limited resource (the robot).

Maja Mataric (University of Southern California) gave an invited talk on action-embedded representations of human-robot interaction. Actions in the form of body language and movement offer a powerful means of communication, not only among people but also between humans and robots. This mode of interaction holds great potential for patient rehabilitation since a robotic "coach" or "nurse" can be used to ameliorate patient engagement without physical contact and, therefore, without the safety and liability roadblocks associated with other forms of rehabilitation robotics.

Panels on assistive technology (chaired by Henry Kautz, University of Washington) and learning from demonstration (chaired by Chad Jenkins, University of Southern California) gave workshop participants a chance to focus on key issues with supervisory control of AI systems. The recurring theme throughout these panels, and many of the technical presentations, was that trust in technology remains a core research issue for future deployment of AI systems. How should the user interface convey to the human operator the capabilities and learning progress of the AI system? Is there a principled way for an AI system to infer the intentions of, and take control from its human supervisor? Other key questions centered around effective strategies and algorithms for machine learning with a human supervisor. Overall, the workshop demonstrated that many opportunities exist for further research at the interface between supervisory control and artificial intelligence.

The papers from the workshop were published as AAAI Technical Report WS-04-10, and are available from AAAI Press.

Michael T. Rosenstein and Mohammad Ghavamzadeh, *University of Massachusetts, Amherst*