

Reports of the AAAI 2011 Fall Symposia

*Sam Blisard, Ted Carmichael, Li Ding, Tim Finin, Wende Frost,
Arthur Graesser, Mirsad Hadzikadic, Lalana Kagal, Geert-Jan M. Kruijff,
Pat Langley, James Lester, Deborah L. McGuinness, Jack Mostow,
Panagiotis Papadakis, Fiora Pirri, Rashmi Prasad,
Svetlana Stoyanchev, Pradeep Varakantham*

■ *The Association for the Advancement of Artificial Intelligence was pleased to present the 2011 Fall Symposium Series, held Friday through Sunday, November 4–6, at the Westin Arlington Gateway in Arlington, Virginia. The titles of the seven symposia are as follows: (1) Advances in Cognitive Systems; (2) Building Representations of Common Ground with Intelligent Agents; (3) Complex Adaptive Systems: Energy, Information, and Intelligence; (4) Multiagent Coordination under Uncertainty; (5) Open Government Knowledge: AI Opportunities and Challenges; (6) Question Generation; and (7) Robot-Human Teamwork in Dynamic Adverse Environments. The highlights of each symposium are presented in this report.*

Advances in Cognitive Systems

The goal of the AAAI Fall Symposium on Advances in Cognitive Systems was to bring together researchers who are interested in developing intelligent systems that demonstrate the full range of human cognitive abilities and to report progress on this daunting task.

The original aims of artificial intelligence, when it was launched in the late 1950s, were to explain intelligence in computational terms and to reproduce the entire range of human cognitive abilities in computational artifacts. Although the field has seen impressive advances in the last few decades, many researchers have, in the process, forgotten or abandoned these important goals. The purpose of the Fall Symposium on Advances in Cognitive Systems was to bring together scientists who remained committed to AI's original vision. The meeting received 50 paper submissions and it was attended by more than 75 participants, suggesting that there remains substantial interest in this view on the discipline.

Research in cognitive systems, as reflected by the contributors to the meeting, differs from what has become mainstream AI in five basic ways. First, it retains a concern with high-level aspects of cognition, such as the ability to engage in multistep infer-



The Advances in Cognitive Systems AAAI Fall Symposium, held at the Westin Hotel in Arlington, Virginia.

Photograph courtesy Will Bridewell.

ence, understand the meaning of natural language, generate novel plans that achieve goals, and even reason about one's own reasoning. Second, cognitive systems research emphasizes the role of structured representations in intelligence and recognizes that both computers and humans are general symbol manipulators, not just numeric processors. Third, it approaches intelligence from a systems perspective, examining how different abilities fit together rather than focusing on component algorithms. Fourth, the field draws ideas and inspiration from findings about human cognition, both for candidate mechanisms and for challenging problems to drive research. Finally, it adopts flexible views on evaluation, welcoming papers that demonstrate new functionality, present novel approaches to established problems, analyze challenging cognitive tasks, and report on architectures for integrated intelligence. Together, these characteristics identify cognitive systems as a separate paradigm that remains committed to the original vision for artificial intelligence.

The symposium included 22 talks and 23 posters

that reflected a number of distinct but connected themes. Some presentations dealt with flexible conceptual-inference mechanisms for using commonsense knowledge that move beyond traditional deductive approaches. A second theme dealt with deep language processing, at both the sentence and dialogue levels, that combines linguistic and domain knowledge to produce genuine representations of meaning. A related topic was reasoning about the mental states of other agents, an ability essential to dialogue and joint activity, with a number of frameworks relying on partitioned memories (divided into "worlds" or "contexts") to keep mental models distinct. A fourth emphasis involved mechanisms for acquiring structural knowledge from experience and instruction, most relying on some form of explanation to generate candidate structures. Together, these themes clarify the distinctive character of research on cognitive systems and its concern with producing human-like computational artifacts.

Open discussions during the meeting focused on three important issues. One concerned criteria for

evaluating research progress, and for publishing papers, on cognitive systems. There was general agreement that existing conferences are not supportive of work in the paradigm, and that the community should find ways to broaden existing venues or establish new ones so that submissions will receive appropriate evaluation. Another discussion topic dealt with challenges that might drive research in the area. Proposals included flexible and scalable inference mechanisms, more natural approaches to acquiring knowledge over extended periods, techniques that interface general-purpose reasoning with specialized methods for spatial and visual processing, and motivational mechanisms that influence reasoning and goal generation. Finally, participants discussed resources that would support the cognitive systems community, with the main candidate being a wiki that incorporates pointers to software, course material, challenge problems, and data sets.

There was general consensus among participants that the symposium served a genuine need within the AI community and that additional meetings were desirable. There was also a sense that, although the cognitive systems paradigm adopts the original aims of AI, its modern incarnation is relatively new and, to ensure its success as a scientific discipline, we must clarify its distinctive characteristics, foster a community of like-minded researchers, identify research challenges and make progress on them, establish venues for publishing results, and train promising new scientists. The symposium was only one step on the road toward the audacious goal of computational systems with the same broad forms of intelligence as humans, and this task will keep the community occupied for years to come.

Pat Langley served as the chair of the meeting. Paul Bello (Office of Naval Research), Nicholas Casimatis (Rensselaer Polytechnic Institute), Kenneth Forbus (Northwestern University), John Laird (University of Michigan, Ann Arbor), and Sergei Nirenburg (University of Maryland, Baltimore County) were coorganizers of the symposium. Papers from the meeting appear in Technical Report FS-11-01 from AAI Press.

Building Representations of Common Ground with Intelligent Agents

The goal of this symposium was to explore methods of explicit or implicit using aspects of common ground to interact with intelligent agents. Much of the success of natural language interaction is caused by the participants' mutual understanding of the circumstances surrounding the communication. The mutual understanding of the perceived context and joint beliefs of the partici-

pants is termed *common ground*, and is made up of all of the background and shared information that will lead to the eventual success of the communication. Some measure of common ground is used in most, if not all, successful interactions between human actors. For humans to have a convincing and beneficial experience interacting with intelligent agents, the agents must have mechanisms that support the fundamentals of common ground. The goal of this symposium was to bring together researchers from diverse fields to examine the representations and models currently in use that will one day enable complete common ground interaction with intelligent agents.

The participants of this symposium came from a variety of AI fields such as robotics, human-computer interaction, linguistics, and cognitive modeling. A major theme of the talks was the need to clarify other agents' beliefs and solidify joint knowledge by asking questions when a situation is ambiguous or unclear. Strategies for doing this and examples of when this might occur were given by Alex Djalali (Stanford University) and Matthew Marge (Carnegie Mellon University). Another major theme of the symposium was the implicit need to build common ground for recognizing and creating deception. This was demonstrated both in cases where the deception is caused by a human, as in a talk presented by Will Bridewell (Stanford Center for Biomedical Informatics Research) during the joint session with the Advances in Cognitive Systems symposium, and in cases when the deception is caused by an intelligent agent, as demonstrated in a paper by Micah Clark (NASA Jet Propulsion Laboratory).

Other topics touched upon by the symposium participants included social authority in dialogue, using outcome matrices for phrase selection, cognitive modeling of common ground, and modeling actual and nonactual states of modal verbs.

The symposium included two invited talks. The first talk was given by Geert-Jan Kruijff (German Research Center for Artificial Intelligence (DFKI)) and was titled "Common Ground in Human-Robot Teams in Urban Search and Rescue." The second talk was given by James Allen (University of Rochester/Institute for Human and Machine Cognition) and was titled "What Are We Doing? A Key Aspect of Common Ground." A joint session was held with the Advances in Cognitive Systems symposium, where topics that were of interest to both symposia were discussed.

The symposium was successful in generating interest among the participants in building the common ground community, which we hope to expand and preserve by future events and additional channels of communication.

Sam Blisard and Wende Frost served as cochairs of this symposium. The papers of the symposium

were published as AAAI Press Technical Report FS-11-02.

Complex Adaptive Systems: Energy, Information, and Intelligence

The goal of the Complex Adaptive Systems symposium was to bring together researchers and students who represent a wide diversity of disciplines, and use the common tools of complex adaptive systems to investigate the next set of fundamental processes in these systems from the perspective of multiple domains.

All complex systems require energy to develop and persist, but “energy” as a concept in complex adaptive systems (CAS) is often more than merely physical energy; it is anything that drives and constrains the system. Agents must cooperate or compete for limited resources, whether these resources are energy, power, food, money, or some other physical or metaphorical limitation. The success or failure of various agent strategies depends on how effective the agents are in acquiring and utilizing these resources.

In this quest for resources, the agent-level implementation of various strategies is enabled by the collection of local information by these agents. But “information,” too, encompasses more than just words, ideas, or memes; it also includes many forms of nonverbal or nonhuman communication systems. We broadly define communication to be any information passing to an agent, from other agents or from its environment. Such flows of information may, in their own right, be the focus of CAS modeling, such as in models of political dissent, social contagion, or the dynamical flows across networks.

Thus, agents in a CAS acquire information from a limited communication radius, interpret what they see, hear, feel, or otherwise sense, and then act based on this knowledge. These actions and reactions spread through a CAS in patterns of correlated feedbacks. This process leads to the emergence of system-level properties, including global patterns of intelligence that no individual agent within that system possesses. Such system-level intelligence can arise even from simple agents, or agents with only a few abilities or limited knowledge. Therefore, these agents cannot be studied in isolation; understanding these complex systems requires understanding how the agents behave and interact with other agents, and in the context of the emergent, system-level properties.

Complex adaptive systems have proven to be a powerful tool for exploring these and other related phenomena. We characterize a general CAS model as having a large number of self-similar agents that: (1) Utilize one or more levels of feedback; (2)

Exhibit emergent properties and self-organization; and (3) Produce nonlinear dynamic behavior. Advances in modeling and computing technology have led to a deeper understanding of complex systems in many areas, and have raised the possibility that similar fundamental principles may be at work across these systems, even though the underlying principles may manifest themselves differently.

In attendance were approximately 45 researchers from disciplines as diverse as computer science, philosophy, economics, political science, biology, public policy, cognitive science, literature, and ecology. Some 20 papers were presented that explored these themes from an equally diverse set of viewpoints. Among the highlights were Kiran Lakkaraju (Sandia National Laboratory), who presented on the diffusion of attitudes in a community; Patrick Grim (Stony Brook University), who presented an analysis of various network structures and how these affect the spread of germs, genes, and memes; and Aaron Bramson, who talked about methods for testing for emergence (University of Michigan). Erika Frydenlund and David Earnest (Old Dominion University) presented two agent-based models: one on community information sharing for improving public safety; and another that explored the resiliency of the international air transportation network. John Seymour and Joseph Tuzo of the University of Maryland, Baltimore County (and their advisor Marie desJardins) discussed two separate models that showed emergent features from the interactions of simple agents: one in the context of lane-changing strategies for highway driving, and the second using an ant colony optimization algorithm. Finally, Russ Abbott (California State University) presented research that compares and contrasts supply-driven ecological systems to demand-driven economic systems.

In addition to paper presentations, we held a participatory panel discussion that highlighted the work from an NEH-sponsored summer institute on simulation for the humanities. Humanists are not usually found at symposia hosted by a computer science organization, and so this was particularly useful in exemplifying how CAS tools and methods can be used in a wide variety of domains. This panel was chaired by Marvin Croy (University of North Carolina, Charlotte) and featured talks from Gillian Crozier (Laurentian University), Graham Sack (Columbia University), and Stephen Crowley (Boise State University).

As with previous years, we were fortunate to have two distinguished keynote speakers. Carl Simon is the founding director of the prestigious Center for the Study of Complex Systems at the University of Michigan. Simon gave us his perspective on the importance of solving practical

problems. We also hosted Sean O'Brien, who is an executive vice-president and senior scientist at SAE, Inc., and a former program director from DARPA. O'Brien discussed various projects and methodologies that affect our armed forces, especially in the context of human, social, cultural, and behavioral (HSCB) dimensions.

Thus, our efforts this year can be summarized as strengthening the community of researchers from across a wide range of fields, and using the three fundamental properties of CAS—energy, information, and intelligence—as a bridge connecting these disciplines. Many of this year's attendees were also present at our first two AAAI CAS symposia, held in 2009 and 2010. Of particular note is the emphasis we place on student participation. Approximately half of the attendees were graduate students or undergrads, including multiple members of the organizing committee.

Mirsad Hadzikadic and Ted Carmichael served as the cochairs for this symposium, and the papers presented here can be found in AAAI Press Technical Report FSS-11-03.

Multiagent Coordination under Uncertainty

The goal of the Multiagent Coordination under Uncertainty symposium was to understand the latest research in multiagent coordination, new application domains, and some of the interesting future directions.

In domains ranging from Earth-observing sensor webs to collaborating ambulances or fire fighters during disaster rescue or software personal assistants scheduling meetings to “coordinators” assisting in executing military missions or exploration of underwater terrains using autonomous underwater vehicles to handling large-scale humanitarian logistics, multiple intelligent agents need to coordinate in the presence of uncertainty to achieve team goals. The goal of the symposium was to understand some of the latest application domains, look at interesting research problems in this area, and investigate future research directions.

The symposium was organized as a set of invited talks by pioneers of the area mixed with talks from researchers who have been working in this area. There were also two discussion sessions where topics raised in the talks and other interesting research questions were discussed. Typically, research in the area has focused on generic models such as DEC-(PO)MDP, DCOP, stochastic games. A key theme that was observed across most of the talks was that the research was based on interesting real-world problems: (1) coordinating security forces to thwart adversaries in airports, country borders; (2) decision support for taxi drivers in Sin-

gapore for better performance (revenue for drivers and higher taxi availability for customers); (3) guiding customers in a theme park for reduced wait times at attractions; (4) building ad hoc robot teams to play soccer; (5) coordinating robot boats to observe interesting phenomenon (observing water temperature and floods to save fish in Philippines); (6) coordinating large-scale logistics services during disaster relief.

There were many exciting and interesting results presented during the talks on latest research. First, in certain large-scale multiagent coordination problems, the quality of solution strategies becomes closer to optimal as the number of agents increases (or the impact of an individual agent on the overall solution is reduced). This was illustrated on route guidance problems faced by cars travelling through a city and for customers at a theme park. Second, selective communication for multiagent coordination can sometimes lead to a fall in coordination performance. Third, multiagent coordination in certain domains with uncertainty can be improved by insights from nature. For instance, robots coordinating for museum surveillance were able to perform better with ant-colony optimization techniques than with standard coordination techniques. Fourth, it is possible to add a new robot to a robot soccer team and perform well without any precoordination between the new robot and the team. Fifth, in certain coordination problems (particularly ones where there is congestion over resources with uncertainty), game-theoretic solution concepts provide better social welfare than locally rational agents (greedy myopic behavior).

There was also considerable discussion on future research directions. An interesting direction raised during talks was about building ad hoc teams of agents in generic problem domains. As part of this, new agents can join teams without any coordination with the team. Reasoning about coordination (and communication for coordination) has typically been performed either completely offline or completely online. Pursuing a hybrid mechanism seems to be a direction for the future. Existing models for coordination under uncertainty (DEC-POMDPs, stochastic games, DCOPs) assume either completely cooperative or adversarial settings and it was agreed that there was need for models that considered partially adversarial and partially cooperative problems. Finally, considering the robust optimization (risk aware) criterion to reason about uncertainty seems to be a natural progression from expected value optimization.

Pradeep Varakantham (author of this report), William Yeoh, Paul Scerri, and Janusz Marecki served as cochairs of this symposium. There were no papers published from this symposium.

Open Government Knowledge: AI Opportunities and Challenges

The AAAI Fall Symposium on Open Government Knowledge focused on issues related to publishing government data as reusable knowledge on the web. More than 40 countries around the world are working on publishing government data as open and reusable world knowledge to improve transparency in government and to accumulate valuable social information. The Open Government Knowledge symposium brought together a diverse community to explore the governmental, business, scientific, and academic challenges and opportunities in the process of publishing, linking, mashing-up, and leveraging public government data.

The symposium featured a wide range of distinguished speakers from government, industry, and academia. Opening keynote talks were given by Jeanne Holm on driving global innovation through linked data and by Abdul Shaikh of NIH/NCI on crowd sourcing of the research enterprise. Invited talks were presented by Debbie Brodt-Giles of NERL on open energy information and by Curt Tilmes of NASA on scientific data and provenance. David Wood of 3roundstones gave a two-hour tutorial on linked data and associated semantic web standards providing the participants with a common understanding of the technology and context. Government practitioners from the United States, Canada, and Brazil reported on their progress and challenges in publishing open data of various kinds including health, energy, community, and education, as well as in preserving the data and its provenance. Speakers from Google (Sree Balakrishnan) and Microsoft (Lewis Shepherd) discussed two important aspects of open government knowledge: web-scale structured data integration and semantic processing to turn data into high-quality knowledge. Several panels addressed advanced topics relating to the role of linked data in an open government knowledge ecosystem, possible “killer apps,” and key technologies for growing business opportunities around open government knowledge.

The symposium generated much discussion among the participants and led to the collaborative development of a road map for linked open government data.

During the discussion, a principle in processing open government data was agreed upon: instead of expecting that all government data be at the highest level of quality, linked data technology can be used to enable incremental data publishing and refinement. Depending on the scale of government organization and the data, the organization might think it is sufficient to get raw data online or may choose to require that all data be directly published as linked data. Another important issue

raised by our government practitioners is the lack of a global, scalable data catalog, which is required to overview the availability of government data and to promote the progressive quality refinement of the data. It was also noted that there are immense quantities of government data that remain in nondigital form. Two important scientific challenges involved with opening reusable and linked government data were identified: (1) management of social-economic issues such as privacy, licenses, access control, risk models, and provenance; and (2) management of computational issues such as uncertainty, correlated metadata and context, special treatment on the temporal-spatial dimension, persistence, and portability.

In summary, open government data has a unique impact on artificial intelligence research. Unlike conventional text web pages, open government data is released and reused in a more structural way, demanding effective AI solutions to enable and enhance cross-domain data integration and web-scale data mash-ups. The huge amount of government data available from different countries, states, and cities further offers a realistic opportunity for evolving computationally and socially scalable AI solutions.

Li Ding, Tim Finin, Lalana Kagal, and Deborah L. McGuinness served as cochairs of this symposium. There were no papers published from this symposium.

Question Generation

Asking questions is a fundamental cognitive process that underlies higher-level cognitive abilities such as comprehension and reasoning. Ultimately, question generation (QG) allows humans, and in many cases artificial intelligence systems, to understand their environment and each other. Research on question generation has a long history in artificial intelligence, psychology, education, and natural language processing. One thread of research has been theoretical, with attempts to understand and specify the triggers and mechanisms underlying question generation. The other thread of research has focused on automated question generation, which has far-reaching applications in intelligent technologies, such as dialogue systems, question-answering systems, web search, intelligent tutoring systems, automated assessment systems, inquiry-based environments, adaptive intelligent agents and game-based learning environments.

The 2011 AAAI Symposium on Question Generation followed three previous workshops and a shared task and evaluation campaign. The goal of the 2011 symposium was to foster theoretical and applied research on computational and cognitive aspects of question generation, thus bringing

together participants from diverse disciplines including natural language processing, artificial intelligence, linguistics, psychology, and education, and to create a forum for discussing and planning the future of question generation shared task and evaluation challenges.

The first part of the symposium featured paper and poster presentations, which represented three major themes. The first theme was human or cognitive perspectives on question generation, under which papers explored the importance of using human models and social interaction for promoting learning, the relationship between curiosity and question asking, and the efficacy of prompts for self-questioning in learning. An invited talk by Patricia Alexander explored the implications of asking questions when prompted or directed to do so by others versus when learners' questions arise spontaneously or intentionally. The second theme focused on various applications of question generation, including generation of multiple choice exams using limited human-judged data resources, generation of questions based on numerical entities, generation of mathematical word problems from ontologies in unrestricted domains, generation of cloze questions based on definitions, generation of questions to promote creativity, generation of questions in situated human-robot interaction, and using automatically generated question-answer pairs for augmenting conversational characters. An invited talk by Jack Mostow reviewed his team's research on automatically generating and administering questions in Project LISTEN's Reading Tutor. The final theme focused on technical approaches for question generation, where research included the generation of more specific questions, ranking of over-generated questions, and exploring the feasibility of using discourse parsing for question-generation tasks.

The second half of the symposium featured break-out working group discussion sessions that addressed evaluation of question generation, specifically through question generation shared task challenge and evaluation campaigns. Evaluating the quality of automatically generated natural language text is notoriously difficult. As natural language is rich and diverse, a correctly generated sentence (in this case, a question) may have multiple forms. This makes automatic evaluation, such as comparing system outputs against a gold-standard key set, difficult. Evaluating natural language generation (NLG) output often involves human judgments of the output in terms of grammatical and semantic correctness, relevance, or appropriateness.

The goal of the discussion sessions was to prepare for the next evaluation challenge. Vasile Rus gave a kickoff invited talk on the first question-generation campaign and future opportunities and

challenges. Participants split into three groups and engaged in a creative discussion on the topics of question-generation evaluation. The organizers defined two tasks for the participants to discuss. Following the discussions, each group presented their ideas and everyone voted on their favorite proposal.

In the first task, participants had to design an enhanced evaluation criteria and process for question generation. Each group had access to the data set and were able to look at the outputs of the systems participating in the first campaign. The groups came up with creative ideas for using crowd-sourcing tools (Mechanical Turk) and enhancing the task to appeal to more participants.

In the second task, participants were asked to brainstorm a different method of evaluation: task-based evaluation. Task-based evaluation measures the impact of an NLG system on how well a user can perform a task while assisted by the system. Task-based evaluation may be an online game, such as navigation in a three-dimensional environment assisted by generation of directions, or a task that participants perform in a lab, such as card matching. Task-based evaluation avoids subjective judgments and, instead, evaluates NLG systems indirectly, by analyzing user performance.

The participants came up with creative new ideas for task-based evaluation. One of the ideas involved an online game where players engage in a Turing test-like task, determining whether a question is generated by a human or by a computer. Correctness of their judgments and quickness of their responses allows evaluation to compare the quality of questions generated by different systems. This type of game may be implemented on popular websites such as Facebook that attract many players and allow for collection of large data sets.

Rashmi Prasad, Svetlana Stoyanchev, Jack Mostow, Arthur Graesser, and James Lester served as cochairs of the symposium. The papers of the symposium were published as AAAI Press Technical Report FS-11-04.

Robot-Human Teamwork in Dynamic Adverse Environments

Robots are gradually making their way into different aspects of our lives. We find them at home, on the factory floor, and, more and more, they are also performing missions in complex outdoor environments.

At the Robot-Human Teamwork in Dynamic Adverse Environments symposium we discussed issues in human-robot teamwork, set in environments that are dynamic and adverse. Typical examples here include urban search and rescue (USAR), or security missions. This is a timely topic: This is already happening. Humans are taking

robots on such missions, using them in situations that are too dangerous for people to (immediately) enter. Where this is not happening yet, it is very likely to be happening soon.

Human-robot teaming in such environments is much more than just a technical issue. It is “people plus robots.” This is a complex sociotechnical system in which humans and robots are trying to work together under very difficult circumstances. They are performing under heavy physical and mental stress, and we know that in such extreme situations, under such extreme conditions, human characteristics, behaviors, emotions are driven to extremes. People change in how they behave, act, what they pay attention to, and how they interact. If we use robots in such situations, will we actually make things better? Or will we make things worse? Can we make robots to assist humans to do better, to handle such situations better?

Looking at these problems from the viewpoint of the user (first responder, police officer, soldier, and so on) the symposium raised from anywhere between practical and fundamental questions. What is a team? What is a human-robot team? What is teamwork in such a context? When we consider mixed-initiative settings, do we actually want robots to take the lead? What makes for good teamwork, and what would be bad teamwork? What does it mean to share situation awareness, when humans and robots are typically geographically distributed, and each and every one has his or her own ways of looking at a situation, experiencing it subjectively? Also, amid all the hows and could, should we?

The take-home messages from the various presentations, and invited talks by Ron Arkin, Jeffrey Bradshaw, and Satoshi Tadokoro, were essentially that it’s about working together. There is an inherent interdependency between humans and robots, and this requires much more than autonomy (even when we consider autonomy to be multidimensional). It’s about the social dynamics of actors, roles, performing particular tasks. It’s about situation awareness specific to interactions between roles, in specific task contexts; and above all, it is about making robots acceptable. As Satoshi Tadokoro aptly put it, it’s perhaps better to turn AI from artificial intelligence into acceptable intelligence: robot intelligence that is clear, predictable, and acceptable in a given team context. Only then we may be able to succeed in turning robots into real team players.

Geert-Jan M. Kruijff, Panos Papadakis and Fiora Pirri served as co-chairs of this symposium. Additional support for this symposium came from the Natural Human-Robot Collaboration in Dynamic Environments project. The papers of the symposium were published as AAAI Press Technical Report FS-11-05.

Sam Blisard is a research scientist in the Interactive Systems section at the Naval Research Laboratory in Washington, D.C.

Ted Carmichael is an assistant research professor in the department of Software and Information Systems at the University of North Carolina at Charlotte, and faculty member in the Complex Systems Institute.

Li Ding is a staff engineer at Qualcomm Inc.

Tim Finin is a professor in the Department of Computer Science and Electrical Engineering at the University of Maryland, Baltimore County.

Wende Frost is a research scientist in the Interactive Systems section at the Naval Research Laboratory in Washington, D.C., and a Ph.D. candidate in linguistics at Arizona State University.

Arthur Graesser is a professor in the Department of Psychology, an adjunct professor in computer science, and codirector of the Institute of Intelligent Systems, University of Memphis.

Mirsad Hadzikadic is a professor in the department of Software and Information Systems at the University of North Carolina at Charlotte, and the founding director of UNC Charlotte’s Complex Systems Institute.

Lalana Kagal is a research scientist in the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology.

Geert-Jan M. Kruijff is a senior researcher at the Language Technology Lab, German Research Center for Artificial Intelligence (DFKI GmbH).

Pat Langley is a professor of computer science and engineering at Arizona State University in Tempe, Arizona, and director of the Institute for the Study of Learning and Expertise in Palo Alto, California.

James Lester is a professor in the Department of Computer Science, North Carolina State University.

Deborah L. McGuinness is a professor of computer science and cognitive science at Rensselaer Polytechnic Institute.

Jack Mostow is a research professor in the School of Computer Science, Carnegie Mellon University.

Panagiotis Papadakis is a postdoctoral researcher at the ALCOR lab (Auto Agent Laboratory for Cognitive Robotics), Department of Informatics and Systems, University of Sapienza, Rome.

Fiora Pirri is a professor of computer science and head of the ALCOR lab, Department of Informatics and Systems, University of Sapienza, Rome.

Rashmi Prasad is an assistant professor at the Department of Health Informatics and Administration, University of Wisconsin–Milwaukee.

Svetlana Stoyanchev is a postdoctoral research scientist at Columbia University.

Pradeep Varakantham is an assistant professor at the School of Information Systems, Singapore Management University.