Reports of the AAAI 2011 Conference Workshops

Noa Agmon, Vikas Agrawal, David W. Aha, Yiannis Aloimonos, Donagh Buckley, Prashant Doshi, Christopher Geib, Floriana Grasso, Nancy Green, Benjamin Johnston, Burt Kaliski, Christopher Kiekintveld, Edith Law, Henry Lieberman, Ole J. Mengshoel, Ted Metzler, Joseph Modayil, Douglas W. Oard, Nilufer Onder, Barry O'Sullivan, Katerina Pastra, Doina Precup, Sowmya Ramachandran, Chris Reed, Sanem Sariel-Talay, Ted Selker, Lokendra Shastri, Satinder Singh, Stephen F. Smith, Siddharth Srivastava, Gita Sukthankar, David C. Uthus, Mary-Anne Williams

■ *The AAAI-11 workshop program was held Sunday and Monday, August 7–18, 2011, at the Hyatt Regency San Francisco in San Francisco, California USA. The AAAI-11 workshop program included 15 workshops covering a wide range of topics in artificial intelligence. The titles of the workshops were Activity Context Representation: Techniques and Languages; Analyzing Microtext; Applied Adversarial Reasoning and Risk Modeling; Artificial Intelligence and Smarter Living: The Conquest of Complexity; Artifiicial Intelligence for Data Center Management and Cloud Computing; Automated Action Planning for Autonomous Mobile Robots; Computational Models of Natural Argument; Generalized Planning; Human Computation; Human-Robot Interaction in Elder Care; Interactive Decision Theory and Game Theory, 2010; Language-Action Tools for Cognitive Artificial Agents: Integrating Vision, Action, and Language; Lifelong Learning from Sensorimotor Experience; Plan, Activity, and Intent Recognition; and Scalable Integration of Analytics and Visualization. This article presents short summaries of those events.*

Activity Context Representation: Techniques and Languages

Pervasive context-aware computing technologies are essential enablers for next-generation applications for the digital workplace, consumer electronics, research, education, government. and health care. These enhanced technologies are expected to be in the mainstream in the next 5–10 years. Context-aware cognitive support requires activity and context information to be captured and, ever more often, moved across devices — securely, efficiently, and with multidevice interoperability.

Task- and activity-based computing originated in the human-computer interaction and computer-supported cooperative work community, such as the ACM Special Interest Group on Computer Human Interaction conference (SIGCHI) and the ACM conference on Computer Supported Cooperative Work. Recently, much work in context modeling, human activity recognition, and modeling using machine-learning techniques and sensor input has been presented in the AI and machine-learning community. As such, activity-based computing lies in the intersection of these two communities. Context representation has received a lot of attention in the information technology community and in the industry among mobile vendors.

This workshop laid the groundwork for techniques to represent activity context using a synthesis of these approaches to reduce demands on people, such as the cognitive load inherent in activity, context, and device switching and enhance human performance within activities. This workshop set the stage for the creation of an international academic and industrial consortium for systems that capture, transfer, and recall activity context across multiple devices and platforms used by people individually and collectively.

The workshop and consortium interest was focused on using AI techniques to improve the human-computer interface for better human performance of knowledge work. Therefore, applications in machine-to-machine systems (manufacturing, smart grid, load balancing), standard data mining, and web-based behavioral analytics were out of scope.

The workshop was introduced by Lokendra Shastri (Infosys) with a vision and motivation for the next-generation digital workspace, where activity context representation techniques will be critical enablers. Jakob Bardram (IT University of Copenhagen) described a powerful implementation of activity-based computing in the health-care space showing examples of benefits from activity context-aware systems. Dan Diaper (author of the *Handbook of Task Analysis*) highlighted that it is far more difficult to determine what context to represent than the representation problem itself, and laid out a formal systems approach for machine capture, representation, and use of context. Arijit Laha (Infosys) described a rich context model for knowledge work. Ontologies to represent context were described by Monica Tentori (Jesus Favela's group on Collaborative Information Retrieval at University of Mexico) and Juan Gomez Romero (University Carlos III of Madrid). Aptima researchers showcased their CHAT (context for human and automation teams) model of context representation.

The lively panel discussion led by Tim Finin (University of Maryland, Baltimore County) defined the scope of the research to include creating context, activity-driven systems providing end-user value through monitoring, exchange and support on activities that can be performed better with help of computational devices than otherwise. Such context includes related information that makes it possible to discern meaning of the information of interest.

The keynote lecture by Henry Kautz (University of Rochester) on activity recognition for way finding and time management illustrated how guidance can be provided by knowing what the user's activity context is. He cautioned that such a smart system could also result in the user having to speculate about what the system might do.

Philippe Palanque (IRIT, Université Paul Sabatier) showed a system providing context-sensitive help for critical system operators. Maarten Sierhuis (PARC), Sonja Zillner (Siemens AG), and Tiffany Tsao (National Taiwan University, Taipei) presented proposals on Brahms, context in medical imaging and a hierarchical activity representation. Fei Li from Schahram Dustdar's Distributed Systems Group (Vienna University of Technology) showed a learning technique for activity recognition.

The second day started with Paul Lukowiz' (Universität Passau) keynote defining outstanding technical challenges in activity and context recognition, capture, representation, and exchange. Aristotelis Hadjakos (Max Muhlhauser's Telecooperation Group at Technische Universität Darmstadt) presented an approach for dynamic context labeling based on product usage. Yasamin Sahaf (Diane Cook's group at Washington State University) showed an example of defining the complexity of an activity. Mobile context-aware systems, context management, and privacy sensitivity were addressed by Boris Moltchanov (Telecom Italia), Tim Finin (University of Maryland, Baltimore County), Tom Lovett (Vodafone), Wolfgang Woerndl (Technische Universität München), Vidya Narayanan and Fuming Shih from Sanjeev Nanda's group at Qualcomm Research and Tim Berners-Lee's group at the Massachusetts Institute of Technology, respectively. Bo Begole (PARC) delivered a closing keynote on the future of activity context-aware systems, technologies, and research.

Based on the strong interest of the participants, the workshop and consortium members will meet at forums such as Intelligent User Interfaces conference (IUI), Where 2.0, ACM CHI 2012, and AAAI 2012 to carry the consortium effort forward. The work includes identifying use-case categories; motivating value with social benefits and business models; creating solution architectures, language, data structures, and operations to enable top use-case categories; significantly augmenting existing standards; creating an adoption plan addressing likely barriers such as critical mass, privacy, not-invented-here, and complexity; providing advisory input to government funding bodies and industry investors; creating fresh initiatives to enable capture, transfer, and recall of activity context and an index of repositories for open-source component software.

Lokendra Shastri, Tim Finin, Henry Kautz, Bo Begole, and Matthai Philipose organized this workshop together with Vikas Agrawal. The organizers thank Gerrit van der Veer (president ACM SIGCHI) for providing publicity among the HCI community for this workshop. The papers from the workshop were published as AAAI Technical Report WS-11-04.

Analyzing Microtext

Text and dialogue analysis is an important area of AI research, and there have been many advances for several genres (for example, news feeds, emails, technical support, and blogs). However, fewer efforts have focused on studying microtext (for example, instant messaging, chat rooms, transcribed speech, and microblogs), which is made up of semistructured pieces of text that are distinguished by their short length, informality, often idiosyncratic lexicon, and sometimes (for example, in group chat) simultaneous interwoven conversations. These characteristics and others (for example, nonstandard grammar, misspelling, and frequent use of emoticons) can make microtext content challenging to analyze.

The 16 presenters at the Analyzing Microtext workshop focused on a broad range of microtext data sources: text chat from internet relay chat and from online games, microblog messages from Twitter and Plurk, transcribed FBI negotiations, Wikipedia barnstars (a type of brief award citation), cell phone SMS (short message service) messages, and image and video captions. Some of the themes covered included normalizing microtext to make it easier to parse with traditional natural language-processing techniques; tagging dialogue and speech acts; topic detection; nontopical content characterization techniques such as sentiment analysis, persuasion detection, and emotion prediction; content-based inference tasks such as threat identification and detection of modified retweets; and applications of machine-learning and related techniques (such as ontology creation).

In addition to the contributed presentations and posters, the workshop included five invited talks from pioneers in microtext analysis. Jeffery Ellen (U.S. Navy Space and Naval Warfare Systems Center, San Diego), the originator of the term *microtext,* described the reasoning that led him to focus on the common elements between these media that have previously been studied mostly in isolation. Abdur Chowdhury (chief scientist of Twitter) described as-yet understudied issues regarding real-time interpersonal information flows in response to major unexpected events such as the Japanese earthquake earlier this year. LorRaine Duffy (U.S. Navy Space and Naval Warfare Systems Center, San Diego) focused on the implications of using tactical chat for mission-critical coordination among geographically distributed and intermittently connected participants. Micha Elsner (University of Edinburgh) stepped back to compare and contrast what we know about spoken dialogue and what we see as emergent behavior in internet relay chat. Craig Martell (U.S. Naval Postgraduate School) completed the picture by presenting examples of cases in which remarkably good results can be obtained from quite simple natural language-processing and machine-learning techniques if they are conceived from the outset with the nature of the medium and the task in mind.

The workshop ended with a discussion panel that was marked by robust and rather enthusiastic audience participation. An important question that emerged concerned future steps. As with many new challenges, it is natural to begin by first evolving existing techniques. There was some consensus that as the research on analyzing microtext matures, the time has come for some researchers to step back and conceptualize the problems anew, and that doing so could help to open up promising new research directions. One challenge in this regard will be to think creatively about how to model these tasks from an evaluation perspective, and then to create corpora that reflect that conceptualization. Another consensus observation was that microtext analysis research presently suffers from fragmentation across too many research communities for any one person to keep abreast. How best to foster interaction among this fragmented community is thus a substantial challenge. There was some support for the idea of developing a wiki to summarize and provide pointers to relevant research (this could serve as a virtual meeting point). There is also substantial interest in additional forums for bringing together researchers from different disciplines (for example, machine learning, computational linguistics, information retrieval, and perhaps social network analysis) to explore potential intersections. For example, we are considering proposing a follow-on meeting as part of the AAAI Symposium Series.

David W. Aha, Douglas W. Oard, Sowmya Ramachandran, and David C. Uthus served as cochairs of this workshop. The papers of this workshop were published as AAAI Press Technical Report WS-11-05.

Applied Adversarial Reasoning and Risk Modeling

Advances in adversarial reasoning have led to exciting new applications, including deployed software assistants for homeland security, poker bots capable of beating expert human players, RoboCup teams with sophisticated adaptive strategies, and tools for managing network and information security. All of these examples share the fundamental challenge of developing agent strategies and decision-making tools that take into account the likely behavior of one or more adversaries. Addressing this challenge in complex real-world domains has inspired many novel tools for adversarial reasoning spanning computational game theory, robust decision making under uncertainty, risk analysis, and opponent modeling.

This workshop brought together researchers

interested in a broad range of different adversarial problems, ranging from poker to homeland security to gold farming in online games. Similarly diverse was the set of methodologies and techniques used to tackle the challenge of adversarial reasoning, including efficient game-theoretic solutions, learning in game playing, crowd modeling, microsimulation, and more. A total of 11 papers were presented at the workshop — 9 full papers and 2 short papers. The papers varied from pure theoretical results in game theory to descriptions of simulation tools to empirical observations of pathologies in some solutions methods.

The workshop included three invited talks, each focusing on a different aspect of adversarial modeling and risk assessment. - Milind Tambe from the University of Southern California described his group's work in different security domains: LAX security, air marshals distribution, and protection of the Boston Harbor. His work represents the security application, and is based mainly on variations of Stackelberg games. - Michael Bowling from the University of Alberta described his work on creating poker-playing agents using a game-theoretic framework. He presented observations on possible pathologies in abstraction approaches, as well as improved methods for calculating worst-case exploitability in very large game trees.

Gal Kaminka from Bar-Ilan University challenged the audience to think more broadly about the scope of adversarial behavior, including behaviors that may be irrational to observers. He presented an axiomatic framework for reasoning about adversarial situations that provides an alternative to the game-theoretic approaches represented in several other talks.

The program and invited talked were very successful at stimulating discussion and debate among the audience, and several common themes among the presented work emerged, including the need to study diverse domains, the need for scalable methods, and the problem of robustness when using incomplete models or abstraction techniques.

Noa Agmon and Christopher Kiekintveld served as cochairs of this workshop, which was also organized by Michael Bowling and Janusz Marecki. The papers of this workshop were published as AAAI Press Technical Report WS-11-06.

Artificial Intelligence and Smarter Living: The Conquest of Complexity

Our world is increasingly instrumented and interconnected. Physical systems and computational systems have almost completely converged: cheap processors are found virtually everywhere and in everything. They are directly touching our lives in everything from toasters, vacuum cleaners, cars, phones, and credit cards; and indirectly through the management of power grids, traffic systems, supply chains, and financial systems. The application of artificial intelligence in this new reality presents an unprecedented opportunity to improve the lives of all people and to create new business value. However, as with any opportunity, there are also significant hazards and risks. Thus, the AAAI 2011 workshop on Artificial Intelligence and Smarter Living was established to serve as a step toward a better understanding of the potential and challenges of artificial intelligence when it is applied to improve lives.

This workshop was intentionally transdisciplinary. We sought to bring together research interests that are not traditionally "close." The workshop attracted researchers not only from artificial intelligence but also from banking, law, and medicine. Papers presented spanned health care, memory assistance, smart homes, smart buildings, banking ecosystems, energy efficiency, medicine, defense, and robotics.

It is clear that smarter living builds upon work in event detection and stream processing, data analytics, sensor networks, and pervasive computing. However, a consistent theme throughout presentations and discussions during the workshop was that smarter living is more than just technology but a complex relationship between users, the environment, robotics, and an intelligent infrastructure. In particular, this is a relationship that is mediated by concerns about privacy, control, independence, and autonomy.

The workshop included a panel discussion that included Ryan Calo (Stanford University Law School), Zico Kolter (Massachusetts Institute of Technology), Daniel Shapiro (Stanford University), and Roland Vogl (CodeX, Stanford University Law School). In addition, two invited talks were given at the workshop. Ben Goertzel (Novamente LLC and Xiamen University) speculated on the requirements and capabilities of future domestic robotics and their software architectures (How much should a robot cost? What sensors does it need? How intelligent does it need to be?). Zico Kolter (Massachusetts Institute of Technology) illustrated the potential of artificial intelligence and machine learning to provide better models of environmental and energy systems that can improve energy efficiency and help users understand their energy impact so they can change their behavior today.

Perhaps the most stimulating discussions arose from the unique combinations of research interests. The workshop revealed a number of possibilities for AI to improve the banking, legal, and health-care ecosystems: first, as a tool to improve the efficiency and sharing among these fields, then later, and more importantly, as an active agent to help unravel (and conquer) the complex interrela-

tionships that stymie the efforts of regulators to improve quality of life.

Benjamin Johnston and Mary-Anne Williams served as cochairs of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-11-07.

Artificial Intelligence for Data Center Management and Cloud Computing

Cloud computing is an emerging paradigm that aims at delivering on-demand computing to any consumer who has access to the Internet. Cloud systems can run software on virtual machines that can be created on-demand in large data centers. These services will be provided through large-scale networks of new data centers, which in turn will connect to the data centers already established by organizations. As a user's demand for computing power increases, new virtual computers can be created and configured. As demand decreases, unused hardware resources can be made available again.

The 451 Group has reported that the use of public cloud computing increased by more than 60 percent in the last quarter of 2008 over the previous two quarters. The International Data Corporation has predicted that spending on cloud services will rise from $16 billion in 2008 to $42 billion by 2012, setting up cloud computing as a key area of growth in the information and communications technology domain. The Gartner Group predicts that more than 80 percent of the top 1000 *Fortune* companies will use cloud-computing services by 2012. Modern data centers are large and, for a variety of reasons related to speed of access, capacity planning and latency, are becoming more and more geographically distributed.

The objective of this workshop was to bring together researchers and technologists from academe and industry to explore the applications of artificial intelligence to the most pertinent technical challenges in data center management and cloud computing. This workshop took advantage of AAAI 2011's proximity to Silicon Valley to attract participants from companies in the area. The workshop successfully attracted the involvement of representatives of companies such as IBM, Google, Microsoft, Hypertable, EMC, and Verisign.

The workshop program comprised two invited talks and several peer-reviewed research papers, giving a full-day program. Doug Judd, chief executive officer and founder of HyperTable, Inc., gave an inspiring presentation on the opportunities for applying AI techniques in cloud databases. Luc Mercier of Google gave an excellent talk about the 2012 ROADEF, Euro Challenge, which focuses on the problem of machine reassignment.

The AI methods that were reported throughout the program were, overwhelmingly, constraint satisfaction, optimization, and machine learning. For example, constraint programming for workload allocation and resource management was a theme in a number of talks. Major challenges in the domains of cloud computing and data center management were related to energy efficiency, cooling, load balancing, and workload allocation. It was noted that the huge scale and multidimensional nature of cloud computing calls for new methods in many subfields of artificial intelligence, where current optimizations have typically been designed for a relatively modest number of objects and a handful of dimensions.

The mix of academic and industry-based attendees was a significant feature of the workshop, which was commented about on several occasions during the event. Industry actively participated in the technical program, with a number of talks being delivered by industry-based speakers. There were many shared interests among the participants, and it was felt that this event could form the basis for an interesting and productive sequence of meetings for people with complementary interests and skills in this important research area.

The workshop cochairs are grateful to all the contributors to the workshop for helping make the event such a success, and to the program committee members who reviewed every submission and selected the papers that were presented Donagh Buckley, Burt Kaliski, and Barry O'Sullivan served as cochairs of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-11-08.

Automated Action Planning for Autonomous Mobile Robots

Robots as embodied agents provide convenient tools to test and validate high-level planners on real-world planning problems. Since the demonstration of Shakey, the first robot that integrated action and motion planning, considerable progress has been made in advancing both action planners and the low-level components of mobile robots that are capable of performing complex planning tasks. However, some difficult challenges still remain.

The Automated Action Planning for Autonomous Mobile Robots workshop aimed to foster the exchange of ideas and to promote research on automated planning representations, models, and algorithms for mobile robots engaged in complex planning tasks. It complements the work in this area presented in recent years at conferences such as AAAI, AAMAS, ICAPS, ICRA, IJCAI, and IROS, and their respective workshop programs.

Eleven papers were presented that explored such themes as planning with missing or uncertain

information, continual task planning, multirobot systems, reconfigurable robots, automatic controller generation, learning spatial information or action durations, and robot-human interfaces. There was an interactive poster session where all participants were invited to present their posters. The presenters of the poster papers were given the opportunity to give spotlight presentations before the poster session.

The first part of the workshop was devoted mainly to automated planning frameworks for mobile robots and dealing with uncertainties while planning. David E. Smith (NASA Ames Research Center) delivered the first invited talk with an interesting title "Planning from Head to Toe." Smith's talk focused on multiple levels of planning for the NASA robot ATHLETE (all-terrain hex-limbed extraterrestrial explorer). Smith also mentioned the main challenges of the automated planning problem for mobile robot systems as temporal features, time constraints, uncertainties, and conflicts. Two different planning and execution monitoring frameworks for mobile robot systems were presented. Some talks presented the use of temporal planners for handling temporal features of the planning problem with a special focus on dealing with failures and uncertainties. Other presentations focused on the use of the partially observable Markov decision process (POMDP) framework for solving partially observable planning problems. Specifically, a hierarchical POMDP approach was presented to solve the overall problem of navigation and visual processing in the face of uncertainty. In another talk, a decision-theoretic planning approach integrated with a classical planner using probabilistic relations was presented as a means to deal with uncertainties.

The second part of the workshop emphasized human-robot interaction and different application domains for planning. The session began with the invited talk "Planning in Personal Robotics" by Bhaskara Marthi (Willow Garage). Marthi outlined three case studies on personnel robotics while emphasizing the main challenges of perception and its interplay with planning. He presented a planning approach for finding hierarchically optimal plans. Marthi ended his talk stressing the need for finding targeted areas where planning fits in robotics. Marthi's talk was followed by two talks that showcased a planning and scheduling framework for human-robot interaction.

The poster session was an integral part of the workshop where participants could interact and discuss different planning approaches for mobile robots. After the poster session, there were presentations for interesting problems including motion planning considering deformable obstacles, reasoning with preferences by situation calculus under erroneous information, and self-reconfiguration. The workshop also hosted an interesting talk on a large-scale scheduling problem for warehouse robots.

Throughout, the workshop was enriched with valuable discussions and several video demonstrations of mobile robots operating in different planning applications. Successful implementations of the current planners on real robots were presented by the representatives from both academia and industry. At the same time, attendees generally agreed that significant challenges remain in the area of automated planning for robot systems, including the interleaving of perception, the handling of uncertainties, and coping with conflicts and failures in real-world execution. Whenever these issues are resolved and planning domains are represented realistically, mobile robots will be capable of performing complex planning tasks in several domains.

Sanem Sariel-Talay, Stephen F. Smith, and Nilufer Onder served as cochairs of the workshop. The papers of the workshop were published as AAAI Press Technical Report WS-11-09.

Computational Models of Natural Argument

The series of international workshops on computational models of natural argument, active since 2001, acts to nurture and provide succor to the ever-growing community working in argument and computation. AI has witnessed a prodigious growth in uses of argumentation throughout many of its subdisciplines: agent system negotiation protocols that demonstrate higher levels of sophistication and robustness; argumentation-based models of evidential relations; group work tools that use argument to structure interaction and debate; computer-based learning tools that exploit monological and dialogical argument structures in designing pedagogic environments; decision support systems that build upon argumentation theoretic models of deliberation to better integrate with human reasoning. This workshop focuses in particular on natural argumentation. Naturalness may involve the use of means, which are more visual than linguistic to illustrate a point, such as graphics or multimedia. It may also involve the use of more sophisticated rhetorical devices, interacting at various layers of abstraction. Finally, naturalness may involve the exploitation of "extrarational" characteristics of the audience, taking into account emotions and affective factors.

Topics of discussion at this and previous workshops have included the characteristics of natural arguments (ontological aspects and cognitive issues); the use of models from informal logic and argumentation theory (in particular, approaches to

specific schools of thought developed in informal logic and argumentation); rhetoric and affect (the role of emotions, personalities, and so on in models of argumentation); the linguistic characteristics of natural argumentation; natural argumentation and media (visual arguments, multimodal arguments, spoken arguments; evaluative arguments and their application in AI systems such as decision support and advice giving); applications of argumentation-based systems; tools for interacting with structures of argument (including visualization tools and interfaces supporting natural, stylized or formal dialogue.); and the building of computational resources such as online corpora related to argumentation.

This workshop featured an invited talk by Olga Gladkova on joint work with Randy Harris and Chrysanne DiMarco on the RhetFig Project (computational rhetoric) and papers by Katarzyna Budzynska and Chris Reed on speech acts of argumentation, by Pranav Anand on annotating persuasive acts in blog text, by Nancy Green on the effect of affect in genetics argumentation, by Ricky Sethi and Yolanda Gil on a social collaborative argumentation system, and by Joel Young on automatic detection of persuasion in microtext corpora. Each presentation was accompanied by lively discussion with the audience. At the end of the day, those participants new to the series expressed enthusiasm for continued participation and all present affirmed how worthwhile the workshop had been to them.

Floriana Grasso, Nancy Green, and Chris Reed served as cochairs of the workshop. The papers of the workshop were published as AAAI Press Technical Report WS-11-10.

Generalized Planning

The second workshop on generalized planning continued the theme of developing new representational methods for increasing the scope and scalability of automated planning, following a similar workshop at ICAPS-2009. Workshop participants presented new ideas and initial approaches for increasing the expressive power of operator representations in planning as well as of the control structure of plans.

The workshop featured three invited talks by prominent researchers in AI. Hector Levesque provided a context for current research on planning with loops with a discussion of the theoretical and applied aspects of computing plans with loops and reasoning about their correctness. Levesque presented the formal foundations of generating plans in the form of finite-state automata (FSAPlans) and also presented some challenging problems in this area for the community.

Stuart Russell's invited talk focused on the value of temporal abstractions captured by hierarchical structures. Russell presented two kinds of hierarchical structures that have been found to improve significantly the scalability of automated planning: hierarchical task networks (HTNs) for nonprobabilistic situations typically addressed in classical planning and partial programs for decision-theoretic planning in the framework of reinforcement learning. Automatically finding such abstractions was highlighted as one of the challenging problems for future research.

Alessandro Cimatti's invited talk described a sophisticated framework for combining automated planning and symbolic model checking for safe and adaptive planning in autonomous spacecraft. Cimatti discussed various aspects of the problem, including the problem of dividing tasks between a control center and the autonomous rover, design validation, adaptive planning, and plan validation. Cimatti's presentation showed a unique integration of symbolic model checking techniques with planning techniques, which is essential in applying planning systems to safety-critical applications.

Paper presentations at the workshop showed the various dimensions along which researchers were generalizing the classical notions of planning, including ideas for generating and analyzing plans with loops for broad applicability; planning with operators of greater expressive power; and for generating plans that are robust to failures in execution and incompleteness in domain models. Toby Hu and Giuseppe DeGiacomo presented two papers developing the notion of plans as finite state controllers. While one of their contributions was a generic solver for constructing such controllers, the second presented new complexity results on generalized planning for finite classes of problems and for one-dimensional planning problems. Ronald Petrick described ongoing work on developing a system for planning in the presence of uncertainty in numerical values of variables. Continuing the theme of planning with loops and numeric uncertainty, Siddharth Srivastava presented a categorization of planning problems with uncertainty in numeric action effects and a sound and complete approach for generating plans with loops for solving such problems.

The postlunch session saw further generalizations of the planning paradigm and included talks on robust planning and planning for service composition as well as on flexible resource scheduling. Christian Muise presented a method for using partially ordered plans (POPs) to build structured policies for execution monitoring as well as for dynamic switching between possible POP linearizations in order to cope with unexpected changes during execution. Marco Pistore presented an approach for service composition while addressing the problems of composition, monitoring, and failure recovery

using ideas from planning and symbolic model checking. Tuan Nguyen presented a measure for assessing the robustness of plans generated with partial models and showed how to compute plans with a given robustness measure by reduction to conformant probabilistic planning. Paul Morris presented a refinement-based approach for producing flexible resource schedules that can deal with execution uncertainty.

The workshop program also included four short presentations and a short open-problem presentation. Workshop participation and presentations of the day reflected a growing, active research community with new contributions, challenging problems for future work, and promising new directions for solving them.

Siddharth Srivastava served as chair of this workshop. Sheila McIlraith, Pado Traverso, and Shlomo Zilberstein coorganized the workshop. The presentations at the workshop were considered preliminary and therefore were not published.

Human Computation

Human computation is a relatively new research area that studies how to build intelligent systems that involve human computers, with each of them performing computation (for example, image classification, translation, and protein folding) that leverages human intelligence, but challenges even the most sophisticated AI algorithms that exist today. With the immense growth of the web, human computation systems can now leverage the abilities of an unprecedented number of Internet users to perform complex computation. Various genres of human computation applications are available today, including games with a purpose (for example, the ESP Game) that generate useful data through gameplay, crowd-sourcing marketplaces (for example, Amazon Mechanical Turk) that coordinate workers to perform tasks for monetary rewards, and identity verification systems (for example, reCAPTCHA) that generate useful data through users performing computation for access to online content.

Over the past few years, we have observed a proliferation of related workshops, new courses, and tutorials, scattered across many conferences. The goal of this workshop is to address this fragmentation by bringing together academic and industry researchers from diverse subfields — machine learning, mechanism and market design, information retrieval, decision-theoretic planning, optimization, human computer interaction — in a stimulating discussion of recent solutions to the core research questions in human computation, and the future directions of this relatively new research area.

The 32 accepted papers and posters fall into three broad themes. The first theme highlights the growing opportunity for AI to function as an enabling technology in human computation systems — to streamline the process of computation, automatically generate workflows, and perform pricing tasks dynamically to meet budget and time constraints. The invited talks by Eric Horvitz ("On Human Computation and Machine Intelligence: Synergies and Frontiers") and Jennifer Wortman Vaughan ("Aggregating Human Predictions Via Markets") further pinpoint the importance of understanding the interplay between human and machine intelligence in computation. The second theme focuses on the "humans" in the loop; several papers present experimental results that reveal new understandings about the motivation, capabilities, limitations, and peculiarity of the human computers, and ways to design systems that can best leverage human intelligence and meet their needs. Finally, there was a wide array of new tools and platforms that address the practical challenges faced by human computation systems in the real world.

This workshop marks the third year that we have organized a workshop on human computation. The number of submissions, as well as their diversity, has grown steadily over the past three years, suggesting the need to continue the human computation (HCOMP) workshop series and eventually move toward expanding it to a symposium or conference format.

Luis von Ahn (Carnegie Mellon University) and Panos Ipeirotis (New York University) cochaired the workshop, along with Edith Law (author of this report) (Carnegie Mellon University), Haoqi Zhang (Harvard University), and Jing Wang (New York University) serving as organizers. The papers of this workshop were published as AAAI Press Technical Report WS-11-11.

Human-Robot Interaction in Elder Care

The Human-Robot Interaction in Elder Care workshop engaged a particularly timely topic. In 2011 the oldest cohort of America's so-called "baby boomer" generation began reaching traditional retirement age — just as factors such as growing health-care costs and shortages of nurses were producing global research and development interest in applications of AI and robotics for elder care.

The wide scope of this interest was reflected in the workshop's attendance, its participants having gathered from Japan, Singapore, New Zealand, Nigeria, and Canada, as well as universities in New Hampshire, Connecticut, Missouri, Oklahoma, and California. Representatives of at least two U.S. federal agencies that fund research in AI and robotics also were present.

Collectively, the participants in this workshop displayed a productive blend of diverse backgrounds and shared research concerns. The gathering was notably multidisciplinary, involving representatives from fields that included engineering, computer science, psychiatry, psychology, gerontology, and philosophy. On the sharing side, though, three of its presenters devoted varying levels of attention to a particular animaloid robot that is used widely in elder care, and two of its presenters both reported using the same humanoid robot in their research. Again, several of the speakers cited research published by other presenters in the workshop.

Such indications of interrelated research efforts notwithstanding, each presentation in the full-day workshop contributed its own unique perspective. The opening speaker reported studies focused on human-robot interaction (HRI) involving older people in their own homes as well as in assisted living facilities, and her report showed careful attention to needs and preferences among human users of the subject technology. The research approach especially attended to measures of improved quality of life for the elder humans who were experimentally interacting with their ostensibly "assistive" robots.

Our speaker from Japan similarly reported research focused upon human factors that can influence the success of HRI, comparing, inter alia, people of younger and older generations. His research made especially clear the potential importance of cultural factors in HRI research. Distinctively Japanese conventions concerning bodily gestures during personal interactions, for example, can become variables of significance in HRI.

Two of the workshop presentations reported research that investigated verbal communication in HRI, but each focused on its own specific aspect of the process. The first report concerned use of human-driven spatial language and involved ongoing joint research by two universities, coupling an electrical and computer engineering department in one with a psychology department in the other. The application scenario of interest in this case was one in which an assistive robot would need to comprehend verbal instructions for, say, fetching an object from another room. The second report concentrated, instead, on technical properties of the human speech signal, formulating a general perceptual model capable of being implemented in a small humanoid robot. One feature of special interest in this model was detection of signal patterns permitting inference of the elder human's emotional state — a feature with evident utility for robotic monitoring of a person's health.

Two other presentations shared interest in ethical dimensions of HRI in elder care. The first of these centrally concerned the animaloid robot mentioned previously. A notable benefit of the presentation was the rich thread of workshop discussion that it produced. Our participants with strong backgrounds in philosophy and in nursing care of elders with dementia appropriately became engaged, leading the dialogue onto topics such as personhood and moral agency.

The dialogue continued in our closing keynote presentation, which reported research on machine ethics in elder care. Following an overview of the prima facie duty approach used in their research, Susan Anderson (University of Connecticut, emerita) and Michael Anderson (Hartford University) gave us a "live" demonstration of their small humanoid robot conducting moral reasoning as it reminded an "elder" (represented by a graphical target) about taking medicine.

The workshop was organized by Lundy Lewis, Susan Barnes, and Ted Metzler (chair). The papers of the workshop were published as AAAI Press Technical Report WS-11-12.

Interactive Decision Theory and Game Theory

This workshop continues a series of workshops on decision and game theories held over several years beginning in 1999. These topics remain active research areas since game and decision theories proved to be powerful tools with which to design autonomous agents, and to understand interactions in systems composed of many such agents.

Decision theory provides a general paradigm for designing agents that can operate in complex uncertain environments, and can act rationally to maximize their preferences. Decision-theoretic models use precise mathematical formalism to define the properties of the agent's environment, the agent's sensory capabilities, the ways in which the agent's actions change the state of the environment, and the agent's goals and preferences. Agent's rationality is defined as behavior that maximizes the expectation of the degree to which the preferences are achieved over time, and the planning problem is identified as a search for the rational, or optimal, plan.

Game theory adds to the decision-theoretic framework the idea of multiple agents interacting within a common environment. It provides ways to specify how agents, separately or jointly, can change the environment and how the resulting changes affect their individual preferences. Building on the assumption that agents are rational and self-interested, game theory uses the notion of Nash equilibrium to design mechanisms and protocols for various forms of interaction and communication, which results in the overall system behaving in a stable manner.

Recent research has sought to merge advances in decision and game theories to build agents that may operate in complex uncertain environments shared with other agents. This research has investigated the problems of Nash equilibrium as a solution concept, focused on epistemological advances in game theory and expressive ways to model agents. Alternative solution concepts have been investigated with the aim of designing autonomous agents that robustly interact with other, highly sophisticated agents in both cooperative and noncooperative settings.

Papers presented at the workshop spanned the spectrum of theoretical issues as well as emerging application areas. There were papers on preference elicitation from a group of users, Bayesian concept learning with application to crowd sourcing, studying the effect of advice given to people, first-order-logic-based epistemological representation, extending Markov games to incomplete information settings, and learning in matrix games. There were also papers on reciprocity in dilemma games, modeling bounded rationality in interactions, and ad hoc teams with application to rovers. The workshop concluded with a substantial discussion on the principle issues confronting the applicability of decision and game theories that emerged from the presented papers, and on ways to broaden the scope of the workshop.

Piotr Gmytrasiewicz, Prashant Doshi (report author), Simon Parsons, and Karl Tuyls served as cochairs of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-11-13.

Language-Action Tools for Cognitive Artificial Agents: Integrating Vision, Action, and Language

The goal of the Language-Action Tools for Cognitive Artificial Agents workshop was to investigate the interface of computational linguistics, computer vision, robotics, neuroscience, and cognitive psychology in developing biologically inspired tools that endow artificial agents with language, action, and perception integration abilities.

Endowing artificial agents with language and action abilities has been a quest in many artificial intelligence subfields. Recent years have witnessed great advances in different disciplines that provide the theoretical and technological framework for an interdisciplinary approach to language-action integration. Neuroscience research provides more and more evidence on a common neural basis for language and action, both in perception and in production. This workshop attempted to shed light on the different aspects of developing biologically inspired language and action technology for artificial agents.

The workshop brought together researchers from a variety of subfields of AI such as computational linguists, computer vision researchers, roboticists, and computational neuroscientists as well as experimental neuroscience and cognitive psychology researchers. Sponsored by AAAI, it comprised 15 invited talks, five paper presentations, and a panel discussion and attracted additional funding from Google Inc. and the EuCognition Network.

One major theme of the invited talks was related to action information as expressed through language; a talk given by Jerry Feldman (University of California, Berkeley) presented simulation semantics as a cognitively motivated computational framework for analyzing verbally expressed events, while Barbara Landau (Johns Hopkins University) focused on comparative data of spatial action language descriptions and the clues revealed by such comparisons regarding basic spatial cognition. Stanley Peters (Stanford University) presented the challenges in grounding even simple verbal references to agents of an event and related work in the analysis of spoken language meeting data. Katerina Pastra (Cognitive Systems Research Institute) presented a generative grammar of action (in the motoric space) and the role of language as a labeling system for action syntax trees, which is generative itself and which requires a new generation of "embodied language technology." Yiannis Aloimonos (University of Maryland) introduced the notion of a cognitive dialogue between perceptual executives, the motor executive, and language executive as a computational framework for developing cognitive artificial agents that are active interpreters.

Another major theme was that of learning language-perception associations. Tamara Berg (State University, New York, Stony Brook) presented work on labeling images and discovering visual attribute terms from web collections. In her turn, Evelyn Tzoukerman (MITRE) presented work on using transcripts and online data for developing joint models of visual action or objects and accompanying text. Ray Mooney (University of Texas) presented language-learning systems that learn from sentences paired with ambiguous, naturally occurring perceptual environments, such as sportscast in simulated RoboCup games and navigation directions in virtual world settings. The use of language as a tutoring tool for learning the meaning of actions that cannot be inferred by observation alone and its use in robot tutoring was elaborated in Britta Wrede's (University of Bielefeld) invited talk. Along similar lines, Gabriela Vigliocco (University College London) advocated the integration of embodied (perceptual, motoric, and affective)

View of San Francisco's Ferry Building from the AAAI-11 Conference Hotel.

experience and information expressed through language in word meaning representations and discussed likely mechanisms for the acquisition of such representations in childhood.

A number of invited talks comprised demonstrations of robots integrating — among other cognitive abilities — perception, action, and natural language; Jeffrey Siskind (Purdue University) demonstrated two robots engaged in playing a board game while a third one observes the play to infer the game rules and a Lincoln logs assembly task undertaken by a robot, while a second one observes that activity and communicates those observations, in natural language, to a third robot who must replicate that assembly. Vadim Tikhanoff (Italian Institute of Technology) demonstrated the POETICON-iCub humanoid succeeding when verbally instructed to perform an everyday task such as "stirring the coffee" without knowing in advance what "coffee" is and deprived of the most commonly used tools for the task.

Finally, neural perspectives of the language-perception-action integration issue were addressed. Tom Dean (Google Inc.) gave an overview of principles that derive from the study of areas of the brain other than neocortex and suggested ways of employing them in computer vision. Max Garagnani (MRC Cambridge) presented a neuroanatomically grounded computational model of sensorimotor circuits for language and action, while Jun Tani (Riken Institute) presented a neurodynamical model of compositionality for goal-directed action generation, mental simulation and planning, free decisions, and language-action associative learning.

The participants discussed the challenges and future directions in this emerging field within a panel discussion moderated by Ruzena Bajcsy (University of California, Berkeley) and Vincent Müller (Anatolia College). Meetings of such a highly interdisciplinary nature were considered by all participants to be the best means for cross-fertilization of ideas, methods, and practices.

Katerina Pastra and Yiannis Aloimonos served as cochairs of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-11-14.

Lifelong Learning from Sensorimotor Experience

A long-standing goal in AI is to understand how an agent can learn from a lifetime of raw sensorimotor experience, autonomously developing internal structures that provide the foundations for further learning and development. Progress is being made, but results are scattered across many subfields with different terminologies and problem formulations,

thus impairing the transfer of relevant ideas between research communities. This workshop brought together researchers with perspectives from robotics, machine learning, and artificial intelligence.

One major theme was understanding specific examples of how a robot can learn from its interaction with the environment. Benjamin Kuipers (University of Michigan) described how a developmental sequence could unfold, where a robot constructs multiple levels of description of its interaction with the environment, including organizing sensors into a spatial model, using spatial models to construct object models, and using object models to formulate high-level actions. Alexander Stoytchev (Iowa State) shared how a robot could learn to push buttons, in particular doorbells, where the auditory feedback provides a clear indicator of success that enables autonomous exploration and the development of reliable behaviors that are adapted to the robot body. Vadim Tikhanoff (Italian Institute of Technology) described aspects of the iCub project, an open robot platform that enables research into several issues in developmental robotics.

Another major theme was understanding how advances in machine learning provide mechanisms that enable lifelong learning in general environments. Richard Sutton (University of Alberta) described how temporal difference algorithms can be used to learn about many facets of a robot's experience. Jürgen Schmidhuber (IDSIA) shared both the early theoretical foundations of self-modifying policies and recent successes enabled by abundant computation. Satinder Singh (University of Michigan) described how optimal reward functions can be formulated to define feedback signals that enable faster learning. Mark Ring (IDSIA) shared research on how an agent's behaviors can be organized in a two-dimensional space enabling generalization. Yoonsuck Choe (Texas A&M University) showed how an agent could learn policies that construct perceptual invariants from edge filters on visual streams.

The workshop also had a lively discussion period, with several ideas from the talks being questioned and clarified. Aspects that were raised included the nature of background knowledge provided for robotics tasks, the resurgence of neural network models, the existence of multiple notions of optimality, and the role of hierarchical structure in a developing agent. The workshop closed with a poster session covering many recent and ongoing projects.

Joseph Modayil, Doina Precup, and Satinder Singh served as cochairs of the workshop. The papers of the workshop were published as AAAI Press Technical Report WS-11-15.

Plan, Activity, and Intent Recognition

Plan recognition, activity recognition, and intent recognition all involve making inferences about other actors from observations of their behavior, that is, their interaction with the environment and with each other. This area of research combines and unifies techniques from computer vision, intelligent user interfaces, human-computer interaction, autonomous and multiagent systems, natural language understanding, and machine learning. It plays a crucial role in a wide variety of applications including assistive technology, security systems, gaming and simulation, proactive assistant agents, human-robot interaction, and multiagent coordination. Several talks introduced improvements on state-of-the-art methods such as lexical plan recognition, using graphical models to analyze sensor data, and weighted abduction. This year, one topic of particular interest was the use of activity and intention recognition within physically grounded AI systems. Chad Jenkins (Brown University) started the workshop with a talk on robotics, presenting an in-depth look at the software behind his group's work on learning by human demonstration. The session also included talks on activity recognition from alternate sensor data, prompting human users within smart homes, and detecting suspicious behavior from human motion trajectories.

Our second invited speaker, Noah Goodman (Stanford University), spoke about social reasoning, the problem of reasoning about other people's beliefs, and presented results on the use of his probabilistic programming language. There was a second presentation on human cognition that focused on autistic planning, and a final session relating to social reasoning in simulated systems such as multiagent teams and coalitions.

The afternoon featured a group discussion about the pros and cons of organizing a plan recognition competition led by Chris Geib (University of Edinburgh) and Chad Jenkins. This session was part of a larger consultation of the community that has included discussions at a Dagsthul symposium on plan recognition and a workshop at ICAPS on plan recognition held earlier in 2011. Although there was enthusiasm about the idea of software and data sets that could result from a competition, there was a general feeling that too simplistic a problem domain could hamper the longer-term research in the area. Participants discussed lessons learned from competitions organized within the planning and robotics communities. There was enthusiasm about gathering existing data sets used in various research papers into a central repository. The workshop concluded with a group dinner, which many of the participants attended.

Gita Sukthankar, Christopher Geib, David Pynadath (ICT), and Hung Bui (SRI) served as cochairs of this workshop. Additional reviewers included Chris Baker (Massachusetts Institute of Technology), Nate Blaylock (IHMC), Aram Galstyan (ISI), Dinh Q. Phung (Curtin University of Technology), and Dorit Zilberbrand. The papers of the workshop were published as AAAI Press Technical Report WS-11-16.

Scalable Integration of Analytics and Visualization

Data and model-rich approaches to analytics are becoming more important across many fields of inquiry — and even in delivering value to customers. The purpose of the workshop was to explore the integration of visualization and analytical techniques, in particular AI techniques, thus enabling computer systems to leverage the abilities of people to perceive and frame problems.

Visual analytics problem solving promises to be a natural, collaborative way of allowing people to drive, debug, and apprehend analytical processes, for example machine-learning results and experiments. In this workshop, we investigated examples of and theory for the way people and their problem solving can benefit from conversationally paced, mixed-initiative visual analytics. One area of focus was interactive exploration of heterogeneous and large-scale data sets, aided by machine-learning, data analysis, information fusion, or statistical techniques.

One question discussed was: What can visualization do for AI? Visualization can help in understanding algorithms and performing visual summarization. In many cases, AI algorithms and systems are black boxes, and visualization may be used to open up the black box. In this way, visualizations can be helpful in fine-tuning parameters of AI algorithms and more generally presenting the complexities of algorithms. AI algorithm animations such as AI space can be useful for educational purposes, and perhaps also increase motivation. Many kinds of semantic spaces that can be visualized were identified: Constraint networks, parse trees, explanations, corpora, time and space, search trees, and Bayesian networks. Anytime there are symbolic structures that are related, one should think about ways to visualize them. If we just made routine the use of visualization techniques whenever anybody in AI presents a semantic representation, we would make the AAAI conference a lot more interesting and accessible to nonexperts.

A second question discussed was what can AI do for visualization? Unfortunately, people may feel overwhelmed by visualization, but perhaps there are ways to generate better visualizations automatically using AI techniques. AI techniques may be used to develop better tools for debugging, better understanding structure, and developing user-adaptive and context-adaptive visualizations. It was suggested that AI could help visualization by making visualization more adaptive to users, more flexible, and take context better into account. Maybe we could use sensors and eye tracking to adjust focus, detail, and others. To make progress in this area we need better understanding of attention patterns, and we need to be mindful of existing research on visualization principles, such as theuse of colors, same or difference relations, and others. Other important issues discussed were scalability and dimensionality; scalability is essential both for scaling up and scaling down (for example, for small screens). Intelligent dimensionality reduction is also necessary. This is why models are useful — they give you the semantics to make those dimensionality reduction decisions. Some people assume there is data and nothing else; however in many cases there is significant background knowledge that should be taken into consideration, including in visualizations.

The lively discussion in this workshop included enthusiasm for an emerging range of ways visualization can help AI practitioners. At one end of the spectrum, we discussed interactive methods for using visualization to help develop and debug Bayesian models. Other systems showed new and productive uses of visualization to help in developing presentation materials and evaluating legislation. It was agreed and demonstrated that AI will commonly be at the heart of creating analytical visualizations. The group found the cross-field inquiry stimulating and expects that it can produce improved user experiences for creating and evaluating knowledge interpretation of every variety.

The participants agreed that the workshop was a success and made progress on the topic of better integrating visualization and analytics. There was clear interest among the participants in attending future workshops with a similar focus as this one.

Ole J. Mengshoel, Ted Selker, and Henry Lieberman served as cochairs of this workshop. The papers of the workshop were published as AAAI Press Technical Report Technical Report WS-11-17.

Noa Agmon is a postdoctoral fellow in the Department of Computer Science at the University of Texas at Austin.

Vikas Agrawal is a research analyst at the Center for Knowledge Driven Information Systems, Infosys Labs, India.

David W. Aha leads the Adaptive Systems Section at the Naval Research Laboratory in Washington, D.C.

Yiannis Aloimonos is a professor in the Department of Computer Science at the University of Maryland, College Park and the director of the Computer Vision Laboratory.

Donagh Buckley is the chief technology officer at EMC's Center of Excellence in Ireland and director, Research Europe at EMC.

Prashant Doshi is an associate professor in the Department of Computer Science at the University of Georgia.

Christopher Geib is a research fellow in the School of Informatics, at the University of Edinburgh.

Floriana Grasso is a lecturer in the Department of Computer Science, University of Liverpool.

Nancy Green is an associate professor in the Department of Computer Science, University of North Carolina Greensboro.

Benjamin Johnston is a senior research associate in the Faculty of Engineering and Information Technology at the University of Technology, Sydney.

Burt Kaliski is senior vice president and chief technology officer at VeriSign, Inc.

Christopher Kiekintveld is an assistant orofessor in the Department of Computer Science at the University of Texas at El Paso.

Edith Law is a Ph.D. candidate in the Machine Learning Department of Carnegie Mellon University.

Henry Lieberman is a research scientist at the Massachusetts Institute of Technology Media Laboratory.

Ole J. Mengshoel is a senior systems scientist at Carnegie Mellon University's Silicon Valley Campus.

Ted Metzler is director of the Darrell W. Hughes Program for Religion and Science Dialogue at Oklahoma City University.

Joseph Modayil is a postdoctoral researcher in the Department of Computing Science at the University of Alberta.

Douglas W. Oard is a professor in the College of Information Studies and the Institute for Advanced Computer Studies at the University of Maryland, College Park.

Nilufer Onder is an associate professor in the Department of Computer Science at Michigan Technological University.

Barry O'Sullivan is a full professor in the Department of Computer Science at University College Cork (Ireland) and director of the Cork Constraint Computation Centre.

Katerina Pastra is the director of the Cognitive Systems Research Insitute, Athens, Greece.

Doina Precup is a professor in the School of Computer Science at McGill University.

Sowmya Ramachandran is a research scientist and project manager at Stottler Henke Associates, Inc.

Chris Reed is a professor in the School of Computing, University of Dundee, UK.

Sanem Sariel-Talay is an assistant professor in the Department of Computer Engineering at Istanbul Technical University.

Ted Selker is an associate director at Carnegie Mellon University's Silicon Valley Campus.

Lokendra Shastri is the general manager, research, and head of Convergence Lab, Infosys Technologies Ltd.

Satinder Singh is a professor in the Department of Electrical Engineering and Computer Science at the University of Michigan at Ann Arbor.

Stephen F. Smith is a research professor in the Robotics Institute at Carnegie Mellon University.

Siddharth Srivastava is a research associate at the University of Wisconsin, Madison.

Gita Sukthankar is a Charles N. Millican faculty fellow and assistant professor in the Department of Electrical Engineering and Computer Science at the University of Central Florida.

David C. Uthus is a National Research Council postdoctoral fellow hosted at the Naval Research Laboratory in Washington, D.C.

Mary-Anne Williams is a professor in the Faculty of Engineering and Information Technology at the University of Technology, Sydney and a fellow in CodeX, the Centre for Legal Informatics, Stanford University.