Reports of the AAAI 2009 Spring Symposia

Jie Bao, Uldis Bojars, Tanzeem Choudhury, Li Ding, Mark Greaves, Ashish Kapoor, Sandy Louchart, Manish Mehta, Bernhard Nebel, Sergei Nirenburg, Tim Oates, David L. Roberts, Antonio Sanfilippo, Nenad Stojanovic, Kristen Stubbs, Andrea L. Thomaz, Katherine Tsui, and Stefan Woelfl

■ *The Association for the Advancement of Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, was pleased to present the 2009 Spring Symposium Series, held Monday through Wednesday, March 23–25, 2009, at Stanford University. The titles of the nine symposia were Agents That Learn from Human Teachers, Benchmarking of Qualitative Spatial and Temporal Reasoning Systems, Experimental Design for Real-World Systems, Human Behavior Modeling, Intelligent Event Processing, Intelligent Narrative Technologies II, Learning by Reading and Learning to Read, Social Semantic Web: Where Web 2.0 Meets Web 3.0, and Technosocial Predictive Analytics. The goal of the Agents That Learn from Human Teachers symposium was to investigate how we can enable software and robotics agents to learn from real-time interaction with an everyday human partner. The aim of the Benchmarking of Qualitative Spatial and Temporal Reasoning Systems symposium was to initiate the development of a problem repository in the field of qualitative spatial and temporal reasoning and identify a graded set of challenges for future midterm and long-term research. The Experimental Design symposium discussed the challenges of evaluating AI systems. The Human Behavior Modeling symposium explored reasoning methods for understanding various aspects of human behavior, especially in the context of designing intelligent systems that interact with humans. The Intelligent Event Processing symposium discussed the need for more AI-based approaches in event processing and defined a kind of research agenda for the field, coined as intelligent complex event processing (iCEP). The Intelligent Narrative Technologies II AAAI symposium discussed innovations, progress, and novel techniques in the research domain. The Learning by Reading and Learning to Read symposium explored two aspects of making natural language texts semantically accessible to, and processable by, machines. The Social Semantic Web symposium focused on the real-world grand challenges in this area. Finally, the Technosocial Predictive Analytics symposium explored new methods for anticipatory analytical thinking that provide decision advantage through the integration of human and physical models.*

Agents That Learn from Human Teachers

Learning will be a key component to the successful application of intelligent agents in real-world environments (both physical and virtual). It will be impossible to give agents all of the knowledge and skills a priori that they will need to serve useful long-term roles in our dynamic world. The ability for everyday users, not experts, to adapt their behavior easily will be key to their success. It is important to recognize that these will be people that are not familiar with machine-learning algorithms; however, everyone has a lifetime of experience with teaching and learning.

Thus, when designing agents that learn from humans, a critical question to ask is "how do people want to teach?" We can draw from social and developmental psychology to find examples of social learning mechanisms (such as imitation and social scaffolding). Then our task is to find computational equivalents for these human-learning mechanisms.

We can also approach the learning from humans problem from the machine-learning perspective. Machine learning is already a process with a "human in the loop," but that human is an expert maching-learning designer. Some examples of the types of tasks that the machine-learning expert has to perform to make an algorithm work are selecting and labeling training data, ordering the examples, deciding when learning is done, defining a reward signal or other optimization criteria. Thus a human is intimately involved in the learning process. Our task is to build frameworks that let nonexperts provide these kinds of

 ISSN 0738-4602

structure and information to machine-learning algorithms.

The Agents That Learn from Human Teachers symposium brought together a multidisciplinary group of researchers for a three-day discussion around the topic of building agents that learn from humans, covering a range of topics. Our three keynote presentations highlighted different perspectives on this problem. From the psychology perspective, we heard from Alison Gopnik about the ways in which young children learn from both their own interventions in the environment and the interventions of social partners. From the human-computer interaction (HCI) perspective, we visited Cliff Nass's lab and heard from a number of his students and alumni about the ways in which people respond to robots and technology. And from the machine-learning perspective, Pieter Abbeel talked about apprenticeship learning as an example of an algorithm that takes advantage of learning from expert human demonstrations. Additionally, several participants presented their work on state-of-the-art examples of software and robotic agents that learn from humans, including five live demonstrations of robot and software interactive learning agents.

Three main themes arose in these discussions: first, design principles for agents that learn from humans; second, algorithms appropriate for this scenario; and third, novel evaluation metrics needed for these types of learning systems.

Andrea L. Thomaz served as chair of this symposium; Sonia Chernova, Dan Grollman, Cynthia Breazeal, Charles Isbell, and Olufisayo Omojokun served as the symposium's organizing committee. The papers of the symposium were published as AAAI Press Technical Report SS-09-01.

Benchmarking of Qualitative Spatial and Temporal Reasoning Systems

Research in the field of qualitative spatial and temporal reasoning (QSTR) is concerned with representation formalisms that are particularly tailored to model continuous aspects of physical reality. QSTR formalisms seem interesting from a cognitive point of view, as they adapt conceptual schemes that are close to human conceptualizations of space and time. On the other hand, qualitative formalisms have attracted sustained research interest within AI because they allow for compact representations of infinite domains and hence promise for efficient procedures to solve spatial and temporal reasoning tasks. In contrast to other established AI communities, however, the idea of benchmarking such formalisms, reasoning methods, and reasoning systems has not played a prominent role so far.

Benchmarking is essential to evaluate new strands of research. This lesson can be learned from other communities in computer science such as automated theorem proving (ATP), Boolean satisfiability (SAT), constraint satisfaction (CSP), and automated planning. The identification of benchmarking problems has had a significant impact on the advancement of these fields. This was also pointed out by Toby Walsh in his keynote lecture, in which he reported on his experiences with the development of CSPLib.org and SATLib.org and on his observations of the CSP and SAT competitions. In a second keynote lecture Geoff Sutcliffe provided a detailed picture of the evolution and the current structure of the TPTP benchmarking library for ATP, which is used in the annual CADE ATP system competition. In a third keynote talk Michael Witbrock identified crucial aspects for building large and broad-coverage knowledge bases such as Cyc. One of the highlights of the symposium was the tool demonstration session, where several participants (partially spontaneously) presented tools and applications resulting from current QSTR research.

The opinion that QSTR could benefit from a problem library was broadly shared among the symposium participants. Further talks and vivid discussion contributions helped to clarify the role of this benchmarking library for future QSTR research topics. In three working groups short- and long-term challenges were identified, and the structure of the problem repository was sketched. Finally, it was suggested to organize a follow-up meeting that should focus on implementation details of the proposed problem library QSTRLib.org.

The organizing committee consisted of Bernhard Nebel (University of Freiburg), Anthony G. Cohn (University of Leeds), Jean-François Condotta (Université d'Artois), Max J. Egenhofer (University of Maine), Ulrich Furbach (University of Koblenz-Landau), Jochen Renz (Australian National University), Peter van Beek (University of Waterloo), Stefan Woelfl (University of Freiburg), and Diedrich Wolter (University of Bremen). The papers of the symposium were published as AAAI Press Technical Report SS-09-02.

Experimental Design for Real-World Systems

As AI systems are deployed into the world, it is important to be able to evaluate their effectiveness. The real world is highly unstructured, and it is important to examine the performance of AI systems in scenarios that test their physical capabilities (for example, a collapsed building after hurricane storm surge has receded) and their social capabilities (such as a crowd of conference attendees at a reception). There are a wide variety of people who will interact with or use these systems, including undergraduate students, search and rescue personnel, geologists and biologists, typically developing children, and children or adults with physical or cognitive disabilities.

The goal of the Experimental Design for Real-World Systems symposium was to discuss the challenges of evaluating AI systems. This symposium brought together researchers from the fields of cognitive science, communication, computer science, electrical engineering, human-computer interaction, human factors, robotics, and social science. Through paper presentations, keynote talks, and panel discussion, a variety of systems were presented, ranging from an ontology describing expert rule systems to a robot arm that classified shadow

puppets made by a person and gestured appropriately in response. The symposium participants discussed how to select appropriate evaluation methods, the selection of type and number of participants, the types and uses of quantitative and qualitative data, and the advantages and disadvantages of various evaluation designs.

The symposium participants debated the appropriateness of types of evaluation methods. For example, the Wizard of Oz (WOz) is a technique used to simulate AI when it is not feasible to provide accurate, real-time interactions. The results from WOz experiments may inform the development and deployment of AI into the real world. An example of a WOz study was shown during a tour of Cliff Nass's Communication between Human and Interactive Media (CHIMe) lab at Stanford, where the symposium participants saw an agent (Robosapien) recommend a movie and then blame the person or itself for an incorrect recommendation.

For systems that incorporate AI techniques, other types of evaluation methods have been used. Marjorie Skubic (University of Missouri) discussed longitudinal observations of a fall-detecting sensor network installed in an aging-in-place facility. Greg Trafton (NRL) showed how cognitive models simulating the gaze behaviors of infants could be evaluated in a parametric study design. Selma Sabanovic (Stanford) demonstrated the use of movie clips in an online survey, a technique often used in psychology research, to determine people's perceptions of gestural interaction. Hatice Kose-Bagci (University of Hertfordshire) demonstrated the use of exploratory studies in child-robot turn-taking while drumming.

The symposium participants joined the Agents That Learn from Human Teachers symposium for a demonstration of embodied and virtual learning agents. The participants discussed how to design evaluations for the demonstrated social learning AI agents and the utility of exploratory and pilot experiments. A variety of performance metrics were found to be applicable to these agents, such as the efficiency of agent learning given limited time for human interaction, the enjoyment of the human when teaching the agent (as determined by self-assessment and psychophysiological sensory information), and the human's assessment of the agent's learning.

The majority of the evaluations discussed were short-term evaluations, which generally occurred in a single session. The longest-running evaluations occurred at the aging-in-place facility and were composed of multiple sessions over the course of three years. As AI systems are introduced into the real world, long-term studies are necessary for AI to achieve the same scientific credibility of other fields, such as anthropology and psychology.

The organizing committee of the symposium consisted of Katherine Tsui (chair) (University of Massachusetts Lowell), David Feil-Seifer (University of Southern California), Heidy Maldonado (Stanford University), Bilge Mutlu (Carnegie Mellon University), Kristen Stubbs (University of Massachusetts Lowell), and Leila Takayama (Willow Garage). The papers of the symposium were published as AAAI Press Technical Report SS-09-03.

Human Behavior Modeling

The participants at the Human Behavior Modeling symposium presented and discussed methods for creating models of individual and group behavior from data. While many researchers across different communities are doing research in behavior modeling, this workshop was distinguished by its emphasis on exploring general representations and reasoning methods that can apply across many different domains. The symposium brought together researchers from a variety of fields including intelligent user interfaces, machine vision, discourse understanding, social network analysis, affective computing, and others. The papers presented in the symposium looked at a wide range of behaviors that include individual activities (such as preparing a meal or taking a walk), interaction between small sets of individuals (for example, having a conversation), and mass behavior of groups (such as the flow of traffic in a city or the collective behavior of millions of cell-phone users). Models included generative and discriminative statistical models, relational models, Markov decision processes, and economic models. Researchers used many different data sources that included GPS, accelerometers, video, speech, cell-phone call logs, and data collected online and through instrumenting users' computers.

The symposium included three invited talks form distinguished researchers active in the area of human behavior modeling. Eric Horvitz from Microsoft Research provided a broad overview of the principles and practices used in human behavior modeling, especially in the context of building mixed initiative applications. Dieter Fox from University of Washington talked about location-based activity recognition and the broad range of modeling techniques his group has developed over the years. On the last day, Alex (Sandy) Pentland from the Massachusetts Institute of Technology talked about honest signals, which are subtle behavioral patterns that reliably reveal people's relationships with others and offer a window into people's financial, cultural, and organizational health. There was also a panel called Resources for Research on Large-Scale Human Behavior Modeling: Opportunities and Challenges. The panelists and participants noted that the lack of publicly available data repositories makes it difficult to compare research done by different groups. However, more and more researchers are making their datasets available to the broader scientific community, which is an encouraging trend.

The participants on the final day discussed challenges in designing scalable techniques, practical issues in experimental evaluation, and a potential killer application of human behavior modeling. Many of the participants shared the focus on developing scalable algorithms that could incorporate multiple different data sources and model individuals as well as group behaviors within a unified framework. Most of the participants also welcomed the idea of organizing and attending symposia with similar focus

in order to advance and share the state of the art in human behavior modeling.

Tanzeem Choudhury, Ashish Kapoor, and Henry Kautz served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-09-04.

Intelligent Event Processing

The vision of responsive business has enforced new challenges for traditional complex event processing (CEP): dealing with unknown events, distributed event processing, and actionable responses, to name but a few. Indeed, not only the scale of processing and fast pattern detection but also the complexity of processing and the efficient discovery and maintenance of complex event patterns determine the success of event processing applications in new, distributed, very complex, and changeable business environments. Most of these new requirements require sophisticated mechanisms like predictability, interoperability, consistency, and completeness, which can be realized by applying formal models and methods. AI seems to be one of the key ingredients for the new generation of CEP approaches, and the Intelligent Event Processing symposium brought together researchers from a variety of subfields of AI to discuss corresponding opportunities and threats. The discussions centered on a number of questions: Why and when intelligent? How intelligent? What is the price of being intelligent? These questions led to investigations about use cases, architectures, and open research issues for intelligent CEP (iCEP). A very inspirational keynote speech by David Luckham (Stanford University) explaining the need for a holistic event processing approach opened the discussions. The main findings of the symposium were as follows.

The primary advantage of iCEP is in dealing with unknown events, which goes beyond the closed-world-assumption-related anomaly detection ("everything that is not known is unknown") to the interestingness-driven unusuality detection ("everything that appears more or less often than expected is unusual"). That means that the system will be able to react to the event patterns that are not defined in advance but generated on the fly according to the background knowledge and past data. This opens opportunities for new applications in the domain of "rare event detection" (such as for crisis management or global or epidemic warning systems).

The key underlying mechanism for iCEP is the logic-based representation of events and operators for combining them in the complex events. The logic provides a unified representation of events, conditions, and actions, enabling reasoning about the reactivity of the system, such as by introducing constraints (synchronization) in the execution of actions. In other words, it is possible to define complex constraints between actions that should be executed as responses on some events, including the case that some of these actions can be treated as new events. Additionally, the logic allows the definition of new operators that describe more complex relations between events than is done by traditional Snoop operators (and, or, seq, and so on) and their efficient realization. An example is the definition of the isolation.

One of the primary research challenges for the community is that the complex event detection process in iCEP must remain data driven. Indeed, logic-based approaches are usually goal driven (backward chaining), which reduces the real-time flavor of event processing.

Another challenge is management of the complex event's patterns, that is, support for all phases in their life cycle, especially creation—a very expensive, SME-driven activity. A pragmatic approach is to treat complex event patterns as knowledge artifacts and apply knowledge-management processes such as creation, representation, usage, and validation. However, automatic methods for complex event mining are even more challenging because iCEP requires the discovery of very complex event patterns and the discovery of new patterns on the fly.

Finally, dealing with complex events on the web is a research challenge on its own, including the questions of an interoperable format for representing events on the web (for example, in RDF) and the corresponding (distributed) complex event detectors.

Nenad Stojanovic, Andreas Abecker (FZI, Germany), Opher Etzion (IBM Research Lab, Haifa, Israel), and Adrian Paschke (Free University Berlin, Germany) served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-09-05.

Intelligent Narrative Technologies II

Narrative is a pervasive aspect of all human societies. Human beings make sense of the world by constructing stories and listening to the stories of others. In addition, stories as a form of entertainment play a central role in our social and leisure lives. As a result, story and narrative have become a key interest for artificial intelligence researchers. The role of narrative as a primary mechanism for organizing human experience has been recognized in many different fields. Work in narrative has become increasingly multidisciplinary with influences from fields including art, psychology, cultural and literary studies, as well as drama. In this context, the symposium focused discussions and presentations on designing computer systems to reason about, perform, and adapt narrative structures for interactive and noninteractive technologies as well as authoring paradigms, tools, and evaluation methodologies for narrative systems.

The symposium brought together researchers from a variety of fields outside of traditional AI such as video gaming, improvisational theater, interactive storytelling, story generation, and story understanding. The symposium had a series of themes and presentations that were loosely organized according to four topics: story generation, social agents, interactive storytelling, and story understanding.

For the story generation session, Reid Swanson from ICT presented a paper he cowrote with Andrew Gor-

don on retrieving information from open domains to generate stories, and Jonathan Rowe introduced StoryEval as an empirical evaluation framework for narrative systems.

In the social agents session, David Knapp talked about using analogies across narratives to drive dialogue and presented a system based on multiagent architecture for which knowledge is coded in English and that features analogical reasoning. Joshua McCoy presented a paper he cowrote with Michael Mateas on the computation of self in everyday life: a dramaturgical approach for socially competent agents.

In the interactive storytelling session, Fox Harrell presented a reflection on agency in interactive storytelling, and Anne Sullivan raised the question of integrating drama management into a playable game experience. Anders Tyschen took the discussion further by suggesting that such issues could be addressed by modeling game master story facilitation in multiplayer roleplaying games.

Finally, in the story understanding session, Beth Cardier discussed the nature of narrative emergence and presented an approach to stories and how we interpret them while James Niehaus explained how humans understand and perform inference with stories and how such knowledge can be applied to computational representations of stories.

There were a number of interesting themes that arose during discussions at the symposium. One of the most interesting came out of a panel discussion on challenges in development and design of interactive narrative authoring systems. Panelists described existing tools for authoring in paradigms other than narrative, the tension between building complete systems and building authoring tools, as well as the set of creative skills that writers use in authoring effective stories. During the ensuing discussion, some participants argued that the performance-oriented nature of stories indicates the need for researchers to abandon tools and focus on building systems. Others argued that the potential disjointness in the sets of skills that technologists and authors have indicates the need for collaborative efforts in building tools.

The AAAI symposia are events that tend to be oriented towards structured discussions, networking, and community building. The Intelligent Narrative Technologies II symposium was no exception as each participant could address the entire audience both through talks (long or short) and posters or demos. The symposium featured two poster sessions, and attendees were exposed to the material of each poster through talks given prior to the poster session. This resulted in very successful poster sessions in which lots of discussions and mutual interests emerged.

Sandy Louchart, Manish Mehta, and David L. Roberts served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-09-06.

Learning by Reading and Learning to Read

The goal of the Learning by Reading and Learning to Read symposium was to stimulate discussion and open exchange of ideas about two aspects of making natural language texts semantically accessible to, and processable by, machines. The first, learning by reading, involves automatically extracting declarative knowledge from text. The second, learning to read, involves automating the process of knowledge extraction required to acquire and expand resources (for example, ontologies and lexicons) that facilitate learning by reading. There is a clear symbiotic relationship between these two aspects — expanding knowledge resources enables systems that extract knowledge from text to improve at that task over time and vice versa.

A distinction that arose early, and that proved to be extremely useful in the discussions that followed, was between macroreading, processing millions of documents (such as web pages) to extract general facts and ontologies, and microreading, processing fragments of a single text to extract deep meaning representations. Tom Mitchell described NELL (never-ending language learning), a system that learns ontologies by macroreading, and noted that the most difficult challenge is extending the time the system can run until it begins to do silly things or to learn things that just don't make sense. This can be done by applying multiple type constraints implied by an ontology, as NELL does, by using redundancy to identify high-confidence facts, or by assuming that highly productive patterns (such as for extracting hypernyms or hyponyms) can be trusted. In all cases the macroreading systems started with seed knowledge, such as an initial ontology or hypernym or hyponym pairs, and then extended this knowledge by processing texts.

Systems that read for deep understanding are potential consumers of the output of macroreaders, but it was not clear to what extent the general ontologies, term taxonomies, and collections of facts they produce are useful for the reasoning required to deeply understand domain-specific text. Curated resources such as WordNet and VerbNet are extremely useful for deep reading, but are not a panacea. One system processed a passage about the attack on Pearl Harbor, found in WordNet that subs, torpedoes, and bombers are all instances of a type of food, and concluded that Pearl Harbor was being attacked by sandwiches. Most such systems presented at the symposium used graph-based representations of text meaning and algorithms that integrated the graphs into larger and larger coherent components. There was general agreement that a large database of domain-specific Schank-style scripts would be a tremendously useful resource for these systems.

The commonalities across the various research efforts were interesting, both for that they did and what they did not do. For example, only one effort, presented by Ken Forbus, learned from an input modality other than text. In this case, Forbus's system learned from text and sketches. While many systems acquired knowledge that could be used to enhance learning by reading, only one, presented by Jon Curtis, truly tackled the problem of learning to read as the primary ob-

jective. Finally, each system used the availability of new text as a trigger for learning, while many fewer used perceived gaps in knowledge (for example, a novel word) to drive focused learning to fill the gap.

Just as learning by reading and learning to read are mutually reinforcing in one system, various research thrusts on macro- and microreading are mutually reinforcing. As the macroreaders produce new knowledge resources, such as collections of inference rules like DIRT, they are used by microreaders to produce better declarative knowledge from the texts they read. The interplay between these two communities, as evidenced at the symposium, can lead to advances in both and a sharpening of approaches to the ultimate shared objective, systems that learn from the vast amounts of text now freely available online.

Sergei Nirenburg and Tim Oates chaired the symposium. The papers of the symposium were published as AAAI Press Technical Report SS-09-07.

Social Semantic Web: Where Web 2.0 Meets Web 3.0

Web 2.0 (social web) applications, such as Wikipedia and FaceBook, are well known for promoting fast growth of online data through network effects. Meanwhile, emerging Web 3.0 applications, driven by semantic web technologies, offer powerful data organization, integration, and query capabilities. The Social Semantic Web symposium investigated joining the strength of data growth power of social web and data organization power of the semantic web to address real-world problems. A special highlight of this symposium was a good number of participants from leading social and semantic web companies such as Vulcan, Yahoo, Oracle, Metaweb (freebase.com), Radar Network (twine.com), Franz, BAE Systems, and InfoSys.

Several interesting themes were observed in this symposium. One focus was on how AI semantic technologies, especially knowledge representation and data mining, could benefit the social web. As data on the social web grows, the need for encoding and organizing complex semantics, such as person identity, community structure, time and location, also grows. Knowledge representaton–based approaches have been presented to encode such complex semantics; moreover, inference problems, such as integrity checking and default reasoning, have been discussed. A number of data mining approaches have been reported to automate the growth of the social web, such as identification of social relation, opinion, community or key tags from social web content, extraction of semantically disambiguated geographic entities and persons in the social web, and network analysis for the online community and tag co-occurrence network.

Another theme focused on how the social web can help the AI community. Tim Finin, a semantic web researcher and professor at the University of Maryland, Baltimore County, gave an invited talk on Wikipedia as an ontology. The talk showed that Wikitology, a hybrid knowledge base constructed from Wikipedia and other knowledge sources, could be used to address many conventional AI and IR problems, such as ontology creation, named entity identification, automated document tagging, and query expansion. Several presentations also suggested that social web–based infrastructures, such as semantic wiki, could curate, evolve, and manage semantic web data, which was previously hard to build and grow.

Yet another theme focused on applications of the social semantic web. Nova Spivak, CEO of Radar Network, gave an invited talk on Twine. The talk demonstrated Twine.com, a semantic social bookmarking and community website for collaboratively filtering web content. It reported interesting statistics about the users of twine (such as high average income and greater than 50 percent of the users contribute or create content), and it also reported lessons learned with challenges in scalability and distributed data computing infrastructure.

One more critical outstanding issue discussed was privacy and trust. A number of presentations covered the legal aspect (for example, privacy laws differ in different countries) and computational aspect (such as access control and policy) of this issue, and many attendants admitted that the semantic web *is* the privacy problem—that is, better semantics may make it easier to reversely derive information protected by privacy laws.

One additional theme was end-user experience. Five panelists from Oracle, Stanford, University of Karlsruhe, Twine, and Freebase and a moderator from Vulcan discussed the requirements and challenges on how to maintain a good end-user experience for social (semantic) web applications. For example, how to best serve the end users without demanding too much learning burden when evolving the underline data organization from tags to RDF or more semantic-intensive structure? How to motivate and facilitate users to collaboratively access, create, evolve, and improve semantic data? How to help connect the social semantic web to our everyday applications (like spreadsheets)? Again, a scalable storage system with low response time is required, and intuitive data access and editing user interface are critical.

Mark Greaves, Li Ding, Jie Bao, and Uldis Bojars served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-09-08.

Technosocial Predictive Analytics

Events occur daily that challenge the security, health, and sustainable growth of our planet and often find the international community unprepared for the catastrophic outcomes. These events involve the interaction of complex processes such as climate change, energy security and reliability, terrorism, nuclear proliferation, natural and human-made disasters, and social, political, and economic vulnerabilities. The goal of the symposium was to foster interactions and partnerships that would help the international community meet the challenges that emerge from these events. The symposium endeavored to achieve this goal through the exploration of new methods for anticipatory analyti-

cal thinking that implement a multiperspective approach to predictive modeling through the integration of human and physical models, leveraging knowledge from both the social and natural sciences, and utilize disciplines capable of supporting the modeling tasks by enhancing cognitive access and facilitating the achievement of knowledge inputs.

The symposium brought together 60 participants, including scientists and government agency representatives, who pondered issues in this emerging field of inquiry with reference to three areas of primary interest: technosocial modeling, knowledge inputs, and cognitive enhancement. The technosocial modeling area targeted the development, implementation, and evaluation of new multiperspective methods and algorithms for predictive modeling. The knowledge inputs area dealt with capabilities that support the modeling task through the acquisition, vetting, and dissemination of expert knowledge and evidence. The cognitive enhancement area focused on the use of visual analytics, enhanced cognition, and gaming techniques to empower the user in the modeling task, promote inferential transparency, and support collaborative and competitive decision making.

The program featured sessions on energy security and reliability, threat modeling and assessment, enabling the user, social and economic simulations, and knowledge management. The two concluding panels discussed current and prospective application domains in technosocial predictive analytics and technical and funding challenges from a government perspective.

The symposium was characterized by intense interaction. Each and every presentation, including those at the panel and poster sessions, generated engaging discussions evidencing the strong interest in this nascent field of inquiry and identifying challenges the community faces in making progress in the relevant areas of focus. We hope that the symposium will stimulate the creation of a new community of interest capable of delivering enduring outcomes by pioneering a new interdisciplinary paradigm of scientific research in proactive critical thinking that can support a concerted decision-making effort by relevant actors to anticipate and counter strategic surprise.

Antonio Sanfilippo served as chair for this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-09-09.

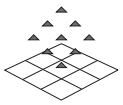

AAAI 2010 Spring Symposium Series

The 2010 Spring Symposium Series will be held March 22–24, 2010 at Stanford University. The call for participation is available in August on the AAAI web site (www.aaai.org/Symposia/Spring/sss10.php). Submissions for the symposia are due on October 2, 2009. Notification of acceptance will be given by November 6, 2009. Material to be included in the working notes of the symposium must be received by January 22, 2010. The complete Call for Participation is available at www.aaai.org/Symposia/Spring/sss10.php. Registration information will be available by December 15, 2009.

Please contact AAAI at sss10@aaai.org with any questions.

Jie Bao is a research associate in the Tetherless World Constellation at Rensselaer Polytechnic Institute.

Uldis Bojars is a member of the the Social Software Unit of DERI Galway.

Tanzeem Choudhury is an assistant professor in the Computer Science Department at Dartmouth College.

Li Ding is a research scientist in the Tetherless World Constellation at Rensselaer Polytechnic Institute.

Mark Greaves is director of knowledge systems at Vulcan, Inc.

Ashish Kapoor is a researcher in the Adaptive Systems and Interaction group at Microsoft Research, Redmond.

Sandy Louchart is a lecturer at the School for Mathematics and Computer Sciences at Heriot-Watt University.

Manish Mehta is a Ph.D. student in the School of Interactive Computing at the Georgia Institute of Technology.

Bernhard Nebel is full professor at the Department of Computer Science, University of Freiburg, Germany.

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

Tim Oates is an assistant professor in the Computer Science and Electrical Engineering Department at the University of Maryland, Baltimore County.

David L. Roberts is a Ph.D. student in the School of Interactive Computing at the Georgia Institute of Technology.

Antonio Sanfilippo is a chief scientist at the Pacific Northwest National Laboratory in Richland, Washington.

Nenad Stojanovic is a project manager and senior researcher at FZI at the University of Karlsruhe.

Kristen Stubbs is the science, technology, engineering, and mathematics outreach program manager for iRobot Corporation.

Andrea L. Thomaz is an assistant professor in the School of Interactive Computing at the Georgia Institute of Technology.

Katherine Tsui is a doctoral candidate at the University of Massachusetts Lowell.

Stefan Woelfl holds a postdoctoral position at the Department of Computer Science, University of Freiburg, Germany.