

Reports of the AAI 2010 Fall Symposia

*Roger Azevedo, Gautam Biswas, Dan Bohus, Ted Carmichael,
Mark A. Finlayson, Mirsad Hadzikadic, Catherine Havasi,
Eric Horvitz, Takayuki Kanda, Oluwasanmi Koyejo, William F. Lawless,
Doug Lenat, Felipe Meneguzzi, Bilge Mutlu, Jean Oh, Roberto Pirrone,
Antoine Raux, Donald A. Sofge, Gita Sukthankar,
Benjamin Van Durme, and Neil Yorke-Smith*

■ *The Association for the Advancement of Artificial Intelligence was pleased to present the 2010 Fall Symposium Series, held Thursday through Saturday, November 11–13, at the Westin Arlington Gateway in Arlington, Virginia. The titles of the eight symposia are as follows: (1) Cognitive and Metacognitive Educational Systems; (2) Commonsense Knowledge; (3) Complex Adaptive Systems: Resilience, Robustness, and Evolvability; (4) Computational Models of Narrative; (5) Dialog with Robots; (6) Manifold Learning and Its Applications; (7) Proactive Assistant Agents ; and (8) Quantum Informatics for Cognitive, Social, and Semantic Processes. The highlights of each symposium are presented in this report.*

Cognitive and Metacognitive Educational Systems

The Cognitive and Metacognitive Educational Systems (MCES) AAI symposium, held in November 2010, was the second edition of this successful AAI symposium. The idea for the symposium stemmed from several theoretical, conceptual, empirical, and applied considerations about the role of metacognition and self-regulation when learning with computer-based learning environments (CBLEs). A related goal was the design and implementation issues associated with metacognitive educational systems. MCES implemented as CBLEs are designed to interact with users and support their learning and decision-making processes. A critical component of good decision making is self-regulation.

The primary aim of this symposium was to continue the discussion started in 2009 on some of the previous considerations and to enhance the discussions with some new ones: What are the theoretical foundations and how are they articulated in CBLEs? Is it possible to develop a unified framework for all metacognitive educational systems? What are the necessary characteristics of these systems to support metacognition? To what extent does the educational system itself have to exhibit metacognitive behaviors, and how are these behaviors organized and enacted to support learning? What are the main aspects of metacognition, self-regulation skills, emotions, and motivations that influence the learning process? What does it mean to be metacognitive, and how can one learn to be metacognitive? Can MCES actually foster learners to be self-regulating agents? How can an MCES be autonomous and increase its knowledge to match the learners' evolving skills and knowledge? What is the role of artificial agents in supporting metacog-

dition and self-regulated learning? MCES may not be embodied, but does it help if they act as intentional agents?

This symposium aimed to provide a comprehensive definition of metacognitive educational systems that is inclusive of the theoretical, architectural, and educational aspects of this field. To meet these goals, we stimulated the debate with two panel sessions. The first, chaired by Janet Kolodner (National Science Foundation) explored the questions of what metacognition is, what pieces of it are needed for learning, what pieces need to be learned, and how can it be supported with technology. The second panel, chaired by Roger Azevedo (McGill University) explored measurement issues in SRL.

The symposium hosted many contributions from researchers in heterogeneous disciplines: AI, cognitive and learning sciences, education psychology, education science, human-computer interaction (HCI), computational linguistics, web technologies, social network analysis, visualization techniques, software architectures, and multiagent systems. Discussion focused mainly on the need to have quantitative measures of the learner's metacognitive abilities. The debate was between education psychologists and AI and HCI people. The former need to have measures of metacognition in support of the evidence of particular behaviors in the learner when he or she is engaged in studying a topic. The latter want to have computable models of metacognitive abilities to build a new generation of truly metacognitive agents that are able to support the learning process. Both kinds of people argued that suitable computable models are needed to represent metacognitive processes despite the particular research goal.

The discussion was enriched by three keynote speeches. Kenneth R. Koedinger (Human Computer Interaction Institute, Carnegie Mellon University) discussed using data to make educational systems more metacognitive. Chris Quintana (School of Education, University of Michigan) gave a speech titled "Making the Implicit Explicit: Issues and Approaches for Scaffolding Metacognitive Activity." Gautam Biswas (Center for Intelligent Systems, Vanderbilt University) presented "Modeling and Measuring Self-Regulated Learning in Teachable Agent Environments." Finally, Janet Kolodner introduced the new National Science Foundation program about cyberlearning and the related funding opportunities.

All the attendees appreciated the theoretical issues that emerged from the discussion of each work and from the invited talks, and they would like to attend future symposia with the same focus as this one.

Roberto Pirrone, Roger Azevedo, and Gautam Biswas served as cochairs of this symposium. The

papers of the symposium were published as AAAI Press Technical Report FS-10-01.

Commonsense Knowledge

When we are confronted with unexpected situations, we deal with them by falling back on our general knowledge or making analogies to other things we know. When software applications fail, on the other hand, they often do so in brittle and unfriendly ways. Applications need commonsense knowledge so that they can interact more sensibly with users and fall back on reasonable defaults, but the sheer amount of commonsense knowledge one would need to represent makes it challenging to acquire and use in applications.

Common sense is, ultimately, the bottleneck to strong AI, and so it has remained one of the central topics of research interest for 50 years, from John McCarthy, Pat Hayes, and colleagues grappling with representation and reasoning, to Doug Lenat, Push Singh, and Lenhart Schubert conducting large-scale engineering projects to construct collections of background knowledge and special-purpose reasoners to support general inference. Recent advances in text mining, crowd sourcing, and professional knowledge engineering efforts have finally led to commonsense knowledge bases of sufficient breadth and depth for practical applications.

A growing number of research projects now seek to use these collections of knowledge in a wide variety of applications, including computer vision, speech processing, robotics, dialogue and text understanding, and apply them to real-world tasks. At the same time, new application domains are giving fresh insights into desiderata for commonsense reasoners and guidance for knowledge collection efforts.

The AAAI Commonsense Knowledge Fall Symposium had the goal of bringing together the diverse elements of this community whose work benefits from or contributes to the representation of general knowledge about the world. We brought together researchers who focus directly on building systems for acquiring or reasoning with commonsense knowledge with those who wish to use these resources to help tackle tasks within their industry or within AI itself.

We observed that commonsense knowledge was united by its goals and applications rather than by a particular formalism. In fact, much discussion took place about the advantages and disadvantages of various formalisms. Two major platforms, Open Mind Common Sense and Cyc, were demonstrated to the symposium and discussed. Several papers were presented that addressed ways to represent common sense and to fill holes in such representations, and we discussed future plans to extend these resources to represent such higher-order

knowledge as causal relationships, justifications, and “scripts” or sequences of events. The symposium served as a place to bring the different theories closer together and to begin to build bridges between the different types of representations used for commonsense knowledge.

The projects represented at the symposium have collected a large quantity of knowledge in several languages through many different methods. We discussed the pros and cons of various methods of knowledge collection: ontological engineering, automatic data mining, and crowd sourcing, including Mechanical Turk-based techniques. We discussed the challenge of how to improve and maintain accuracy of knowledge bases in light of such rapid knowledge expansion and proposed solutions from the point of view of various formalisms.

We discussed the role crowd sourcing has played in knowledge acquisition, especially using the games with a purpose (GWAP) model. Participants presented many innovative games, including the PTT Virtual Pets game, the Turing Game, and Cyc’s FACTory, that provide input to commonsense knowledge bases. In the other direction, we discussed ways that common sense could help in game AI, such as Ken Forbus’s (Northwestern) system, which uses common sense to form high-level strategies in the game FreeCiv.

The combination of this large amount of knowledge with easy-to-use resources, large-scale data-mining knowledge-collection projects, and new models of reasoning, indicates that we have reached an inflection point in our field. Common sense has evolved to the point where others are easily able to use our resources in their systems and applications. Many applications of common sense were presented including word-sense disambiguation, opinion mining, recommender systems, color selection, literature search aids, textual entailment, medical informatics, and cognitive modeling. We talked about supporting commonsense tools in deployed applications, such as Ken Forbus’s work with sketch understanding tools and Catherine Havasi’s work with the Media Lab’s Glass Infrastructure project.

Representatives from many different project initiatives attended the symposium as well as researchers who used commonsense knowledge in their applications ranging from story creation to game design to sentiment analysis. The symposium featured many productive discussions and was very useful in spurring increased collaboration in the field. Participants felt the symposium increased understanding and collaboration within the field and seemed eager to attend another symposium on this topic.

Catherine Havasi, Doug Lenat, and Benjamin Van Durme served as cochairs of this symposium.

The papers of the symposium were published as AAAI Press Technical Report FS-10-02.

Complex Adaptive Systems: Resilience, Robustness, and Evolvability

Companies, societies, markets, and humans rarely stay in a stable, predictable state for long. Yet all these systems are characterized by the notable persistence of some key attributes that maintain their identities, even as their constituent parts change and adapt to new environments. What is it about these systems that define their identity? How do we characterize them? What are the forces that allow a system to persist, even in the face of a radically new environment? Complex adaptive systems (CAS) have proven to be a powerful tool for exploring these and other related phenomena. We characterize a general CAS model as having a large number of self-similar agents that: (1) utilize one or more levels of feedback; (2) exhibit emergent properties and self-organization; and (3) produce nonlinear dynamic behavior. Advances in modeling and computing technology have led to a deeper understanding of complex systems in many areas and have raised the possibility that similar fundamental principles may be at work across these systems, even though the underlying principles may manifest themselves differently.

For some practitioners in the field, the terms *resilience* and *robustness* may seem largely redundant. Indeed, there are many other terms from various domains that overlap as well: from *basins of attractions* (physics, mathematics), to *homeostasis* (biology), to *sustainability* (ecology). This is precisely the point: different disciplines often have their own language, even as they are describing identical or similar phenomena.

In attendance were approximately 35 researchers from disciplines as diverse as computer science, philosophy, economics, political science, biology, and cognitive science. Papers were presented that explored these themes from an equally diverse set of viewpoints. Among the highlights were Matthieu Branlat and David Wood’s perspectives on resilience engineering; Anthony Beaver’s work analyzing evolving dynamic networks; Mary Rorick and Gunter Wagner’s presentation on the adaptive evolution of proteins; Bob Reynolds, Xiandong Che, and Mostafa Ali’s work on problem solving with cultural algorithms; and Kiran Lakkaraju and Ann Speed’s paper on cognitive models of populationwide attitude changes.

In addition, the symposium included a participatory panel on the common language of CAS. Panelists Chris Eichelberger, Aaron Bramson, Matthieu Branlat, and Ted Carmichael each presented a short perspective; and discussant Mirsad

Hadzikadic led a very robust discussion with all panelists and participants. The theme of this panel was elucidating various approaches for finding common grounds and a common grammar of CAS, so that cross-disciplinary work might truly be interactive across scientific boundaries. This is a continuing challenge for this emerging field.

The symposium concluded with a half-day tutorial on agent-based modeling, led by Bill Rand and Forrest Stonedahl, and two invited speakers who gave talks on their current research. The invited speakers were Patrick Grim, distinguished professor of philosophy from Stony Brook University, and Maggie Eppstein, professor of computer science at the University of Vermont and director of the Complex Systems Center.

Thus, this symposium can be summarized as strengthening the community of researchers from across a wide range of fields who use the key phenomena that characterize system resilience as a bridge across these disciplines. Indeed, many of this year's attendees were also present at the first AAAI CAS Symposium in the fall of 2009. Of particular importance is the emphasis placed on student participation, resulting in a symposium where over half of the attendees were graduate students and undergrads—and even a high school student—all of whom were interested in undertaking CAS research. This symposium was supported in part by the National Science Foundation grant #1052901.

Mirsad Hadzikadic and Ted Carmichael served as the cochairs for this symposium, and the papers have been published in AAAI Press Technical Report SS-10-03.

Computational Models of Narrative

The three main goals of the 2010 AAAI Symposium on Computational Models of Narrative were to survey the state of the art regarding the fundamentals of the computational representation and modeling of narrative, to identify gaps and next steps, and to help build a broad, interdisciplinary community. The symposium built upon the momentum generated at the 2009 Work-

shop on Computational Models of Narrative.¹

By the measures of the organizers and the participants, the symposium was a great success. A main contributor to that success was the wide range of papers and viewpoints and that the event was truly international and interdisciplinary. Nearly 40 researchers from over 10 countries in the Americas, Europe, and Asia attended the symposium, and they brought expertise from many areas, including commonsense reasoning, formal logic, natural language processing, language generation, representational formalisms, analogical reasoning, legal reasoning, argumentation theory, geospatial narratives, interactive narrative technologies, cognitive science, cognitive narratology, linguistics, discourse analysis, cognitive psychology, anthropology, sociology, and philosophy. The breadth of scholarship present expanded the view of all participants and greatly enriched the experience.

The symposium began with Roger Schank's keynote address, in which he gave an account of his recent whereabouts working on collecting and using stories in the corporate world. He gave many fascinating examples of how narratives from one domain (for example, transoceanic shipping or corporate negotiation) can inform another (for example, drug development). He set the tenor of the symposium by demonstrating both technical precision and narratological sophistication.

As noted, the symposium followed on the 2009 workshop with the expectation that we would push forward the envelope of our understanding. While there were numerous new ideas, perhaps the most notable was the advance of our view of the range of complexity of narratives. It was clear that most computationalists focus on quite simple narratives, that perhaps wouldn't even be called such by humanists. Finlayson presented a review of research, prepared especially for the symposium, which showed that the vast fraction of computational approaches to narrative used but a single story, usually less than three sentences long!

In contrast, the humanists attending the symposium gave the computationalists a taste of much more complex

narratives. Barbara Dancygier, a cognitive narratologist from the University of British Columbia, described a scheme for visualizing changes of viewpoint in a narrative.² As an example, she analyzed a portion of the novel *The Blind Assassin*, and showed how there were no fewer than six narrative viewpoint switches. This clearly presents a challenge for computationalists, whose simple stories invariably contain only a single viewpoint.

Similarly, Loren Niemi, a professional storyteller, gave a performance and discussion of the art of storytelling that highlighted a range of different types of narratives. He sketched for us numerous versions of *Little Red Riding Hood*, as it has been told in different historical eras. What was interesting was how, although remarkably different in surface form, these narratives all told the "same" story. Understanding and measuring this sameness, which exists at a level far above the types of representations computationalists have so far considered, presents a whole new vista.

Charlotte Linde, a sociologist from the NASA Ames Research Center described her work on collecting and studying how stories are used to create value or meaning relative to social membership in a group.³ One fascinating point she described was how the same story could mean different things to different people in different contexts, and that to even discuss the "meaning" of a story, one must take into account numerous widening circles of context and culture. In the discussion it became clear that these circles of context were far above the level of representation at which the computationalists were used to working.

These examples of complex narratives set the stage for a productive discussion about constructing a story corpus to facilitate and focus work in the area. At last year's workshop it was agreed that a "story bank" corpus could do for computational narrative what the Penn Treebank has done for computational linguistics, namely, provide a common substrate to save work and time, and on which to make reasonable comparisons. Much of the symposium's discussion on the morning of the third day centered on constructing

such a corpus. Of great concern to the participants was that during the initial stages the corpus should not focus too strongly on a single genre, modality, or representational scheme, the thought being that different researchers have quite different needs. One valuable idea that emerged was to begin by constructing a seed corpus containing a “handful of handfuls,” that is to say, a small set of narratives from a small, broad set of genres, say, 10 narratives from each of 10 genres spanning many modalities and topics, for example, folktales, short stories, novels, plays, movies, comics, business school cases, medical discharge summaries, legal briefs, historical notes, and so forth.

The remainder of the discussion focused on what direction to take the meeting series—when and at what venue should the next meeting take place? Should the next meeting range broadly like the symposium and workshop, or focus more on a specific problem, such as constructing a corpus? What features should the next meeting have to maintain the breadth of scholarship, international participation, energy, and momentum present at last year’s workshop and this year’s symposium? While many excellent suggestions were fielded, too numerous to mention here, it is notable that the participants were unanimous that the series should continue. This is perhaps the clearest and most convincing evidence that the symposium was a success, and we look forward to the next meeting.

Mark Finlayson served as chair of the symposium. The papers of the symposium were published as AAAI Press Technical Report FS-10-04, edited by M. A. Finlayson, P. Gervás, E. Mueller, S. Narayanan, and P. H. Winston.

Dialog with Robots

Researchers in the human-robot interaction (HRI) community have addressed a spectrum of challenges at the intersection of robotics, social and cognitive psychology, human factors, and AI. At the same time, progress is being made in the spoken dialog community including work on the development of fundamental theories, formalisms, and models. To date, spoken dialog efforts have focused almost exclusively on ap-

plications within restricted communication contexts, such as telephone- and PC-based information access. Several research efforts at the intersection of spoken dialog and HRI have sought to broaden dialog to richer, more natural settings. These efforts have identified numerous challenges with the use of dialog as part of coordination among multiple actors, taking into consideration details of the tasks at hand and the surrounding environment.

The AAAI Fall Symposium on Dialog with Robots was organized to catalyze communication and innovation at the crossroads of spoken dialog and HRI. The meeting brought together over 70 researchers from the HRI, spoken dialog systems, intelligent virtual agents, and other related research communities in an open discussion about the challenges at the intersection of these fields. The proceedings contain over 40 contributions.

Ideas spanning a spectrum of inter-related research topics were presented and discussed during oral presentations and a poster session. Recurrent themes centered around challenges and directions with the use of dialog by physically embodied agents, taking into consideration aspects of the task, surrounding environment, and broader context. Several presentations highlighted problems with modeling communicative competencies that are fundamental in creating, maintaining and organizing interactions in physical space, such as engagement, turn-taking, joint attention, and verbal and nonverbal communicative skills. Other presenters explored the challenges of leveraging physical context in various language understanding problems such as reference resolution, or the challenges of coupling action and communication in the interaction planning and dialog management process. A number of papers reported on developmental approaches for acquiring knowledge through interaction and focused on challenges such as learning new words, concepts, meanings, and intents, and grounding this learning in the interaction with the physical world. The topics covered also included interaction design challenges, descriptions of existing or planned systems,

research platforms and toolkits, theoretical models, and experimental results.

Invited keynote presentations provided different perspectives on work in dialog with robots and complemented views and themes arising at the symposium. Cynthia Breazeal outlined progress in recent years in the HRI community, including developments in social robotics and uses of models of mind in robotic systems. Candace Sidner provided a review of the state of the art on modeling discourse and dialog in the spoken dialog community and highlighted challenges in pushing the boundaries of the current models at different interaction time-scales, from seconds to hours to months. Herbert H. Clark added a valuable psycholinguistics perspective to the meeting by sharing his reflections about how people coordinate joint activities. He presented results from studies with human subjects on the timing of utterances during collaborations, highlighting sets of competencies required for fluid spoken interaction and collaboration.

In addition to the technical and keynote presentations, the symposium included three moderated, open discussions that provided a forum for exchanging ideas on some of the key topics in the symposium. The first discussion aimed to address some of the challenges at the crossroads of dialog and HRI. The physicality of such interactions was highlighted as a critical factor, and the prospect of identifying a core, yet simple set of principles and first-order concepts to be reasoned about, or a “naïve physics” of situated dialog and discourse, was raised and discussed. A second open discussion centered on the interplay between action and communication and highlighted ideas such as viewing communication as joint action and the importance of creating models for incremental processing that can support recognition and generation of actions and phenomena occurring on different time scales. The final discussion addressed several other fundamental issues such as how we might move forward in this nascent field. Discussion touched on the need for unified platforms and challenges for supporting comparative evaluations of dif-

ferent techniques, the pros and cons of simulation-based approaches, and even the value of revisiting fundamental questions: Why should we endow robots with the ability to engage in dialogue with people? What assumptions are we making—and which can we make?

Altogether, the density of ideas captured in the technical contributions and participants, the animated discussions, and the diversity of questions raised, and the different technical approaches taken reflect a nascent, vibrant community unified by the common themes of dialog with robots or interaction within a physical context. We thank the authors and participants for their valuable contributions and perspectives on these problems and hope that the momentum generated by this symposium will carry forward and help to catalyze future research efforts and meetings.

Dan Bohus, Eric Horvitz, Takayuki Kanda, Bilge Mutlu, and Antoine Raux served as cochairs of this AAAI symposium. The papers of the symposium are published as AAAI Press Technical Report FS-10-05.

Manifold Learning and Its Applications

Many modern problems involve the analysis of high-dimensional data sets. Researchers have observed that in many cases, these high-dimensional data samples have some significant structure; particularly in local neighborhoods. The goal of manifold learning research is to discover techniques that exploit local structure to learn more effective or efficient models for describing high-dimensional data. Manifold learning has become a truly cross-disciplinary field, involving researchers from various subject areas such as topology, geometry, machine learning, statistics, computer vision, robotics and many others. This has led to an accelerating pace of research and applications in recent years. The goal of the symposium was to provide a forum for discussing the current state of the art in theory and applications and to discuss common open problems.

The symposium was organized into five sessions with six invited speakers,

presentations of submitted papers, and moderated discussions. The ideas discussed ranged from spectral methods, to probabilistic models, to multimani-fold models and data dependent dictionary learning. Yi Ma (University of Illinois at Urbana-Champaign) presented new computational tools for extracting rich low-rank structures in images and videos. These models have achieved state-of-the-art performance for many tasks in computer vision and image processing such as feature extraction, image alignment, three-dimensional reconstruction, and object recognition. Neil Lawrence (University of Sheffield) presented a probabilistic perspective on spectral dimensionality reduction, showing how several common spectral methods such as Isometric mapping (ISOMAP), local linear embedding (LLE), and maximum variance unfolding (MVU) can each be described as special cases in a common probabilistic framework based on Gaussian random fields. He then showed the performance of the generalized model for robot navigation and for describing motion capture data.

Fei Sha (University of Southern California) discussed a large class of unsupervised kernel dimensionality reduction algorithms. He also proposed several models for supervised dimensionality reduction where the goal is to simultaneously reduce the dimensionality of the input and learn an appropriate predictor for the output. He showed how such models can lead to superior performance and to more interpretable results using examples in weather prediction and handwritten digit classification. Lawrence Carin (Duke University) presented a class of nonparametric Bayesian models for manifold learning. He showed how the trained model described a chart for the manifold and could reliably estimate many important manifold parameters such as the intrinsic dimension. He also discussed several applications of these techniques including a time-evolving model for predicting votes in the United States House of Representatives.

Gilad Lerman (University of Minnesota) discussed the foundations of multimanifold data modeling where the data is generated by many, possibly

intersecting manifolds. He showed several theoretical and practical results, particularly in the case of linear subspaces, and showed that the resulting algorithms are very robust to noise and outliers. He also showed several applications to motion segmentation and feature tracking in video data. Guangliang Chen (Duke University) presented new ideas for data-dependent dictionary learning, where one learns a multiscale dictionary for describing high-dimensional data sets. He showed how new samples can be encoded using the dictionary and discussed the performance guarantees.

The symposium included extensive discussions on common open problems. For instance, some theoretical aspects of manifold learning are not yet well understood. These include constraints on the types of manifolds that can be learned and the number of data samples that are required for learning a reliable estimate. Participants suggested the development of a more general framework for manifold learning encompassing many of the current techniques as special cases. Such generality would simplify the description of many current models, clarify modeling assumptions, and may suggest more effective techniques. Participants discussed the need for improved model validation. Many manifold learning algorithms are designed for visualization and are validated using heuristic criteria. The participants proposed that we solve this problem by using benchmark data sets and more objective criteria so that different algorithms can be easily compared. Objective criteria for choosing the best-performing models will also ease the adoption of these techniques by other practitioners. Participants discussed the danger of testing manifold learning algorithms on toy data sets like the “swiss roll,” as the intuition gained may not be applicable to real noisy and undersampled high-dimensional data. Finally, though there are many compelling applications, participants discussed the need for a “killer app” where the manifold assumption is clearly the most appropriate and effective. As our understanding continues to grow, we expect that many of these issues will be resolved and manifold learning will continue to

be an important tool for high-dimensional data analysis.

Oluwasanmi Koyejo authored this report. Oluwasanmi Koyejo and Richard Souvenir served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-10-06

Proactive Assistant Agents

From its inception, agent research has espoused the creation of systems capable of providing assistance to human users in everyday tasks such as making travel arrangements for business trips or vacations, helping children with their homework, and reminding forgetful users of important information. Thus, the pursuit of ideal assistants has been persistent, leading to vigorous research in intelligent assistant agents in various problem domains ranging from personal assistants to military applications. However, many challenges still remain before widespread adoption of proactive assistant agents is a reality. These challenges are inherently multidisciplinary, and only by joining expertise from the breadth of the AI community can these challenges be addressed.

To facilitate a wide-ranging discussion of the unique technical challenges arising in intelligent assistant systems, this symposium was intended to be a forum for integrating and advancing the state of the art in agent-based technologies from various fields. As a result, we had participants working on assistant agents from diverse groups including academia and both industrial and military research labs. The specific research areas included agent technologies, intention recognition, planning and scheduling, robotics, goal reasoning, social network analysis, knowledge capture, and cognitive science.

To nourish active discussion, the symposium was formatted to emphasize vigorous debate about each individual paper. Paper authors were instructed to limit their presentation to the minimum needed to seed the discussion, with most of the session time spent discussing the work itself with all participants. This special format led to lively and engaged discussions dealing with not only the state of the art but al-

so key issues and potential avenues for further research in proactive assistants.

Specific topics of presented papers included a variety of domains. The first session was devoted to assistant agents in social media, with applications in news reading and social networking. The second paper session dealt with various applications of proactive assistants, such as formal modeling of tutoring systems, assistance for data handling in scientific workflows, as well as planning using multiple sensor data to provide effective assistance for the elderly. The final paper session was focused on intention recognition, including papers on novel techniques to infer human intention, as well as the application of intention recognition for the particular problem of elderly care.

One of the specific goals of this symposium was to gather the researchers from various projects in assistant agents to share their wisdom in retrospect. Invited talks included former members of various mature assistant agent projects such as Electric Elves (celebrating its tenth anniversary), CALO, and RADAR. The invited speakers were then gathered in a panel session focused on the lessons learned from those implemented systems. During the panel session, one of the recurrent issues was that researchers need to focus on the usability of the prototype assistants at an early stage of development. More of the issues discussed in the retrospective session will be amalgamated in a journal paper currently in progress.

Trends on military applications of assistant agents were also the subject of invited talks, with presentations on research in the Army and Navy research labs. Two key issues were discussed in these talks. One was the need to introduce proactive assistance methods to help personnel from multiple organizations (both military and nonmilitary) reconcile policy conflicts. The other was the need to develop goal-directed autonomous reasoning to allow assistant agents to really be able to adapt the assistance provided to unforeseen circumstances.

An interesting focal point raised and discussed repeatedly was the intrinsically multidisciplinary nature of assis-

tant technology and the need for embracing more collaboration with human-computer interaction researchers. The take-home message is that the main challenges of proactive assistants remain as difficult AI problems; at the same time, careful attention must also be paid in designing how an assistant interacts with people: after all, attitude matters.

Felipe Meneguzzi and Jean Oh served as cochairs of the symposium. This report was authored by Felipe Meneguzzi, Jean Oh, Gita Sukthankar, and Neil Yorke-Smith. The papers of the symposium were published as AAAI Press Technical Report FS-10-07.

Quantum Informatics for Cognitive, Social, and Semantic Processes

Quantum informatics (QI) is the science, technology, and application of quantum information, including its use in abstract models of macroscale systems and processes. Researchers from a variety of disciplines including artificial intelligence, cognitive and brain sciences, decision theory, logic, data storage and retrieval, social science and organization theory, economics, finance, and ecology have recently begun exploring the use of QI to address challenging problems in their respective fields where classical and traditional methods have proved inadequate.

In the fall of 2010 researchers from these disciplines (and others) met to share knowledge and research results on applying QI to cognitive, social, and semantic processes. This was the fourth in a series of meetings that began with the first quantum informatics symposium held at Stanford University in March 2007 as part of the AAAI Spring Symposium Series. This was followed by successful meetings in Oxford, UK, in 2008, and Saarbrücken, Germany, in 2009. The focus of these symposia is the application of QI to model systems outside the domain of quantum physics and to gain understanding and appreciation for how mathematical methods, representations, and models that have found great success in characterizing quantum phenomena in physics may be applied to understand-

ing and characterizing macroscale complex systems such as language, cognition, and social interactions in groups.

The symposium was organized by themes including semantic processing, conceptual processing, dynamics and systems, quantum cognition, foundations, and information retrieval. Keynote presentations were offered by Pentti Kanerva (Stanford) on prototypes and mapping in concept space and by Terry Bollinger (MITRE) on a quantum-inspired approach to interpreting sensor data. In addition, there were panel-led group discussions including the topics “What Makes a System Quantum” and “Grant Opportunities in QI-Space.”

The session on semantic processing included a presentation on evaluation of a computational model of abductive reasoning using predication-based semantic indexing and a novel and efficient approach to representing words and concepts as vectors in high-dimensional space. The session on conceptual processing included a presentation on the separability of ambiguous combinations of words and concepts in joint probability space, a presentation on modeling entrepreneurial decision-making processes in the context of external environmental factors, and a presentation on nonclassical effects found in modeling combinations of words found on the World Wide Web, known as the Guppy effect.

The dynamics and systems session included a presentation on self-regulation of organizations based upon the conservation of information and a new model of interdependence and a presentation on a QI model of the dynamics of interacting species. This was followed by the quantum cognition session with presentations on a new QI model of episodic memory in human cognition and a cognitive heuristic (based on a quantum formalism) used to provide a promising new way to model human judgment and categorization of concepts. Also presented was a quantum formalism that successfully describes human strategic decision making in certain game-theoretic contexts where classical Bayesian methods and Nash equilibria fail.

The foundations session included a

proposed computational logic-based grammar for constructing self-organizing systems, and an analysis of how quantum concepts are often misunderstood due to our propensity to construct incorrect pragmatic interpretations of them. The final session on information retrieval explored the application of QI to the storage and retrieval of documents. It included presentations on using tensor products to store correlated text and image features to facilitate image retrieval and on QI models of users’ relevance judgment and cognitive interference in comparing two documents.

Peter Bruza, William Lawless, Dominic Widdows, and Donald Sofge cochaired the symposium, while Kirsty Kitto and Jerome Busemeyer served as session chairs. Keith van Rijsbergen served on the Organizing Committee but unfortunately could not attend. This report was authored by Donald Sofge and William Lawless. The papers of the symposium were published as AAAI Press Technical Report FS-10-08.

Notes

1. See M. A. Finlayson, W. Richards, and P. H. Winston, 2010. Computational Models of Narrative: Review of a Workshop. *AI Magazine* 31(2): 97-100.
2. See B. Dancygier, 2010. The Language of Stories: A Conceptual Integration Approach. In *Computational Models of Narrative: Papers from the AAAI Fall Symposium (AAAI Technical Report FS-10-04)*, pp. 14-15. Menlo Park, CA: AAAI Press.
3. See C. Linde, 2010. Social Issues in the Understanding of Narrative. In *Computational Models of Narrative: Papers from the AAAI Fall Symposium (AAAI Technical Report FS-10-04)*, pp. 39-40. Menlo Park, CA: AAAI Press.

Roger Azevedo is full professor in the Department of Educational and Counseling Psychology, McGill University, Canada.

Gautam Biswas is full professor in the Center for Intelligent Systems, Vanderbilt University.

Dan Bohus is a researcher at Microsoft Research, Redmond.

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

Mark Finlayson is a doctoral candidate in

the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology.

Mirsad Hadzikadic is an associate professor in the department of Software and Information Systems at the University of North Carolina at Charlotte and the director of the Complex Systems Institute.

Catherine Havasi is a postdoc at the Media Lab, Massachusetts Institute of Technology.

Eric Horvitz is a distinguished scientist at Microsoft Research, Redmond.

Takayuki Kanda is a senior research scientist at ATR Intelligent Robotics and Communication Laboratories in Kyoto, Japan.

Oluwasanmi Koyejo is a Ph.D. student in the department of Electrical and Computer Engineering at the University of Texas at Austin.

William Lawless is a professor in the Departments of Mathematics and Psychology at Paine College in Augusta, Georgia.

Doug Lenat is president and founder of Cycorp in Austin, Texas.

Felipe Meneguzzi is a research fellow at the Robotics Institute at Carnegie Mellon University.

Bilge Mutlu is an assistant professor at the Department of Computer Sciences at the University of Wisconsin, Madison.

Jean Oh is a research fellow at the Robotics Institute at Carnegie Mellon University.

Roberto Pirrone is associate professor in the Department of Computer Science and Engineering (DINFO), University of Palermo, Italy.

Antoine Raux is a scientist at the Honda Research Institute USA.

Donald Sofge is a computer scientist in the Navy Center for Applied Research in Artificial Intelligence at the Naval Research Laboratory in Washington D.C.

Gita Sukthankar is an assistant professor at the Department of Electrical Engineering and Computer Science at University of Central Florida.

Benjamin Van Durme is a research scientist in the Department of Computer Science at Johns Hopkins University.

Neil Yorke-Smith is an assistant professor of business information and decision systems at the American University of Beirut, Lebanon, and a research scientist at SRI International’s Artificial Intelligence Center.