

AAAI 2006 Spring Symposium Reports

Andreas Abecker, Rachid Alami, Chitta Baral, Tim Bickmore, Ed Durfee, Terry Fong, Mehmet H. Göker, Nancy Green, Mark Liberman, Christian Lebiere, James H. Martin, Gregoris Mentzas, Dave Musliner, Nicolas Nicolov, Illah Nourbakhsh, Franco Salvetti, Daniel Shapiro, Debbie Schrekenghost, Amit Sheth, Ljiljana Stojanovic, Vytas SunSpiral, and Robert Wray

■ The American Association for Artificial Intelligence, in cooperation with Stanford University's Computer Science Department, was pleased to present its 2006 Spring Symposium Series held March 27–29, 2006, at Stanford University, California. The titles of the eight symposia were (1) Argumentation for Consumers of Health Care (chaired by Nancy Green); (2) Between a Rock and a Hard Place: Cognitive Science Principles Meet AI Hard Problems (chaired by Christian Lebiere); (3) Computational Approaches to Analyzing Weblogs (chaired by Nicolas Nicolov); (4) Distributed Plan and Schedule Management (chaired by Ed Durfee); (5) Formalizing and Compiling Background Knowledge and Its Applications to Knowledge Representation and Question Answering (chaired by Chitta Baral); (6) Semantic Web Meets e-Government (chaired by Ljiljana Stojanovic); (7) To Boldly Go Where No Human-Robot Team Has Gone Before (chaired by Terry Fong); and (8) What Went Wrong and Why: Lessons from AI Research and Applications (chaired by Dan Shapiro).

Argumentation for Consumers of Health Care

Different notions of argument historically have played a central role in artificial intelligence to model the diagnostic reasoning and decision making of medical experts. However, it was beyond the scope of that research to address information needs of the layperson. It was assumed that a medical expert, trained to interpret explanations produced by a system, would mediate between system and layperson. The goal of this symposium was to investigate the role of argumentation in future intelligent health-care systems, especially focusing on systems designed to interact directly with health-care consumers.

The symposium brought together researchers from a variety of subfields of AI such as decision support, multi-agent systems, computational models of natural argument, embodied conversational agents, and natural language text generation, in addition to researchers in fields providing empiri-

cal or theoretical foundations (including health communication, argumentation, social psychology, cognitive psychology, psycholinguistics, and human factors). One major theme of papers presented at the symposium was health behavior change support systems, interactive systems designed to encourage users in activities such as healthy eating and exercise. (These papers continued a major theme of the AAAI 2004 fall symposium on health dialog systems.) The symposium also included two invited talks on this theme. A talk given by Stacy Marsella (University of Southern California / Information Sciences Institute) focused on the use of behavioral theory and emotion modeling in design of pedagogical dialogue systems. A talk given by B. J. Fogg (Stanford University) outlined general principles of persuasion in human-computer interaction.

Another major theme was normative argument in health care, from decision-support systems for medical experts to systems generating patient-tailored explanations designed for the health-care consumer. The papers on this theme described systems using a variety of models of argument (such as from argument theorists such as Toulmin and Walton). Several papers provided analyses of argument patterns found in actual medical discourse. Although the use of argument in such systems has a different purpose than in health behavior change systems, the symposium participants discussed the potential need to consider a patient's emotional state in presenting normative arguments to patients. The third theme of the symposium was linguistic and cognitive factors in effective communication of medical information. These papers served to remind the participants that the language, layout, and graphics in which an argument is conveyed can affect the audience's comprehension of the argument and that presentation factors can even be manipulated to prevent an audience from accurately assessing negative information.

The symposium participants discussed how interest in argumentation is growing but is still fragmented among many fields and how the sym-

posium was useful in bringing together many different perspectives of argument. However, the participants shared the goal of developing methods to improve health outcomes and agreed that they would like to attend future symposia with the same focus as this one.

Timothy Bickmore and Nancy Green served as coauthors of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-06-01.

Between a Rock and a Hard Place: Cognitive Science Principles Meet AI Hard Problems

Herbert A. Simon once wrote that "AI can have two purposes. One is to use the power of computers to augment human thinking.... The other is to use a computer's artificial intelligence to understand how humans think." Today, the notion that AI and cognitive science have common goals is debatable. AI as practiced today often focuses on defining and measuring domain-specific approaches to problems. References to human approaches to the problem are the exception, not the rule. Cognitive science has largely merged with cognitive psychology: evaluation and progress are usually measured in this field by rigorous empirical research on tiny slices of human behavior. The "20 questions" approach to psychology still dominates psychology, 30 years after Allen Newell cautioned against it. The purpose of this symposium was to bring together researchers to discuss the accuracy and implications of Simon's notion that AI and cognitive science share the same goals.

The symposium was organized around the presentations of papers that responded to one or more of seven questions outlined in the call for papers: (1) Is cognitive science relevant to AI problems? (2) Are "good enough" solutions valuable? (3) Are multilevel heterogeneous approaches beneficial? (4) Is adaptiveness an essential component of intelligence? (5) Are the most efficient solutions prob-

lem specific? (6) Is developing specialized, modular components a reasonable way to study general intelligence? (7) Can artificial intelligence contribute to our understanding of human cognition? The papers ranged from nontechnical position papers to technical results highlighting a perspective on a question. One of the more novel and exciting aspects of the symposium was the evidence of the growing influence of brain and developmental sciences on our shared understanding of "intelligence."

Perhaps unsurprisingly, the majority of the submissions to the symposium were largely in agreement with Simon's view. Three speakers were invited to present less commonly shared views to the participants. Richard Korf (UCLA) started the symposium with a presentation on why AI and cognitive science may have little in common, which resulted in a lively start to the symposium. David Kieras (University of Michigan) argued that the lessons of modern computer organization and operating systems can be used to update the "mind as information processing" model developed to inform psychologists about likely cognitive processes and requirements. Pat Langley (Stanford University) advocated taking a unified approach to intelligent systems constrained by cognitive science principles, in contrast to the typical software-engineering-based integration approaches of AI. The invited speakers all did an excellent job of challenging assumptions and provoking passionate discussion.

This symposium continued an ongoing discussion around the desirability and feasibility of the general notion of "human-level intelligence," most recently at the AAAI 2004 fall symposium chaired by Nick Cassimatis and Patrick Winston. In our closing session, we discussed potential solutions to the barriers facing researchers in this area, including the high cost of entry, immature methodologies and evaluation criteria, results sharing across different paradigms, and software interoperability. The specific actions taken include establishing a mailing list and website¹ to help foster a more cohesive community around the goal of human-level intelligence.

Christian Lebiere and Robert Wray served as coauthors of this symposium. The papers of the symposium were published as AAAI Technical Report SS-06-02.

Computational Approaches to Analyzing Weblogs

Weblogs are web pages that provide unedited, highly opinionated personal commentary. Often a weblog (also referred to as blog) is a chronological sequence of entries that include hyperlinks to other resources. A blog is conveniently maintained and published with authoring tools. The blogosphere as a whole can be exploited for outreach opinion formation, maintaining online communities, supporting knowledge management within large global collaborative environments, and monitoring reactions to public events, and it is seen as an upcoming alternative to the mass media.

The goal of this symposium was to bring together researchers from different subject areas (including computer science, linguistics, psychology, statistics, sociology, multimedia and semantic web technologies) interested in using AI tools for understanding the blogosphere and to foster discussions about ongoing research in the following areas and topics: AI methods; blogosphere versus mediasphere; centrality or influence of bloggers; ranking and relevance of blogs; crawling and indexing; human computer interaction; multimedia; semantic analysis; semantic web; sentiment analysis; social network analysis; text categorization; time-series forecasting; trend identification.

The invited talk was presented by Geoffrey Nunberg who talked about blogging in public. Another significant session was the industrial panel titled Technologies to Understand the Blogosphere Now and Gain Insight in the Future, which featured representatives from major industry players Umbria, Feedster, Six Apart, Yahoo, AlwaysOn Network, Cymfony, and Google. Interesting questions from colleagues from Nielsen and Buzzmetrics stirred the discussion.

The symposium attracted presentations on the structure of the (Ameri-

can) blogosphere; identifying subcommunities; recommendation of other bloggers with overlapping interests; opinion, sentiment, and mood identification, the happiness of the bloggers during the day and the week; spike analysis (given an increased volume of blogs over a short period of time what are the main events that caused the increased activity); the demographic analysis of the bloggers (based on bloggers' profiles and directly inferring gender from the writing style); summarization of salient information; studies of the cultural differences among bloggers in terms of observable characteristics like volume, frequency of blogging, predictions models from blogosphere activity to sales (such as movies, proposals for multilingual processing), and so on. A critical issue that the symposium brought to the foreground is that the blogosphere has been heavily polluted with spam blogs (splogs), which undermines the results of any models if the problem is not addressed head on.


In the spirit of fostering a community, we created a weblog² where participants posted interesting comments during the sessions.

The feedback to the symposium was very positive; in the words of one participant: ". . . not the typical NLP-mix of topics, which I actually enjoyed quite a bit—seeing new perspectives on problems, and the audience was a great mix of academia and industry."

The papers from this symposium were published in the AAAI technical report series and are available from AAAI press. A selection of extended and revised papers will be published by the MIT Press. We will follow up with an international conference on weblogs and social media in Boulder in 2007.

"And remember that if you are blogger, we can discover your gender, your age, your interests, where you live, what your happiness level is, and who is going to win the next election."

Nicolas Nicolov, Franco Salvetti, Mark Liberman, and James H. Martin served as cochairs of this symposium. The papers of the symposium were published as AAAI Technical Report SS-06-03.



The Spring Symposia Were Held at Stanford University in the History Corner.

Distributed Plan and Schedule Management

Automated systems for planning and scheduling the activities of coordinated groups must scale into increasingly complex, dynamic, stochastic, and even adversarial application domains where plan or schedule failure can be costly or catastrophic. In such domains, distributed plans and schedules must be managed on an ongoing basis to avoid, tolerate, correct for, or even exploit evolving circumstances.

This symposium brought together researchers tackling various aspects of the distributed plan and schedule-management problem to foster a broader understanding of the problem, to establish a fruitful dialogue between people who have been focusing on insular pieces of

the problem, and to initiate an exchange of ideas, techniques, and best practices.

To promote discussion, each symposium session was organized around broad and important questions. For example, several sessions revolved around the question of how to formulate and solve aspects of the distributed plan and schedule problem using techniques such as distributed constraint reasoning, distributed Markov decision processes, distributed resource allocation mechanisms, and distributed temporal constraint network representations. Other sessions focused on issues of managing the computation and communication required by distributed plan management, including issues of coordinating the timing of management activities across agents, handling changes and outages in the communi-

cation infrastructure, and making trade-offs between the quality and timeliness of distributed management decisions.

The distributed plan and schedule management problem at the heart of the Defense Advanced Research Projects Agency (DARPA) Coordinators program served as a frame of reference for much of the discussion. A significant fraction of participants were involved in Coordinators efforts. Thus, once the problem of coordinating the activities of military units in a volatile setting was discussed and dissected, it helped create a vocabulary and context that facilitated further comparison and discussion of techniques beyond those currently being funded by DARPA. Moreover, a variety of other application domains such as robotics, space systems, industrial design, interactive games, and logistics also helped enumerate challenge problems for distributed plan and schedule management.

Whether for commercial, military, or other uses, a common theme in the deployment of these technologies in real-world settings was the difficulties in gaining user acceptance. Symposium participants shared their experiences in introducing automation for supporting peoples' plan and schedule-management activities and identified key concerns in human-computer interfaces, in balancing stability in peoples' schedules with systemwide flexibility to respond to emergent needs, in modeling the criteria that people use in making management decisions, and in engaging human expertise only at appropriate times.

While there are periodic symposia and workshops on constituent theories and technologies related to distributed plan and schedule management, this symposium brought together researchers to talk about the broader issues that must be tackled to solve problems in automating distributed plan and schedule management in complex, realistic contexts. At the symposium's conclusion, it became clear that many were interested in continuing the conversation.

Edmund Durfee and David Musliner served as cochairs of this symposium. The papers presented at this symposium are available as AAAI Technical Report SS-06-04.

Formalizing and Compiling Background Knowledge and Its Applications to Knowledge Representation and Question Answering

This symposium topic was motivated by the observations that most knowledge representation languages are flat and currently there is no openly available literature on developing knowledge bases in a modular manner. This is in contrast to the development of programs in languages such as C and Java; one needs to develop large knowledge bases for applications such as question answering; and to develop large knowledge bases, one needs a way to be able to reuse modules developed earlier and/or by others.

The symposium consisted of three invited talks, several oral presentations, and a panel. The three invited talks were by Richard Fikes, Peter Clark, and Doug Lenat. Fikes gave a background on earlier funded and ongoing efforts for developing large knowledge bases and discussed lessons that were learned from those efforts. Clark talked about WordNet, text mining, and knowledge bases of the future. Doug Lenat talked about CYC and explained how important subsets of CYC are now available to outsiders.

A big part of the oral presentations focused on various proposals for modular knowledge representation languages. Lifschitz and Ren presented their work on modular action languages; Baral, Anwar, and Dzifcak presented their work on modular declarative logic programming and illustrated it with respect to reasoning about actions and planning; Gelfond presented a slightly different modularization and illustrated it with respect to a knowledge base that reasons about trips; Khatib, Pontelli, and Son presented their work on integrating Prolog and answer set programming; Gelfond and Zhu presented their work on combining logical and probabilistic reasoning; and Gabaldon presented his work on hierarchical task libraries in Congolog. Besides the above there were two presentations on specific

knowledge base development using modularization techniques: one by Scherl and another by Lee and Lifschitz. A completely different approach of using semantic skeletons for building a repository of background knowledge was proposed by Galitsky.

Besides Doug Lenat's invited talk on CYC, there was one more presentation on the syntax and content of CYC based on a paper by Matuszek, Cabral, Witbrock, and DeOliveira.

Ontologies as well as context, two important issues in building knowledge repositories were explored by McCarthy and McIlraith in two separate presentations. While McCarthy talked about approximate objects, McIlraith presented her work with Grunigner on specifying a web service ontology in first-order logic.

The panel had two parts: first, technical issues for developing knowledge repositories, and second, the next steps in making the idea of a knowledge repository a reality.

Leora Morgenstern remotely organized the first part by posing questions such as: What are some of the technical problems that researchers will need to solve in order to develop sizable knowledge repositories?

The second part of the panel was led by Chitta Baral. He steered the discussion towards making a list of action items that will lead to the building of large knowledge repositories. One of the action items was that interchange standards should take into account knowledge representation languages beyond first-order logic, especially nonmonotonic logics, and answer set programming.

Overall the symposium was a great success, and Michael Gelfond represented the symposium in the plenary session. Chitta Baral served as chair of this symposium. The papers presented at this symposium are available as AAAI Technical Report SS-06-05.

Semantic Web Meets E-Government


The main intention of this symposium was to bring together two communities: the semantic web and the e-government researchers. On one hand, af-

ter several years of a very successful research, the semantic web community is still searching for real use cases that will profit from the machine-readable or understandable description of information. On the other hand, e-government research is searching for innovative technologies that can cope with the ever demanding challenges in automation of public services, like cross-national interoperability.

In fact, this symposium served as the kickoff meeting for the semantic e-government community, a community at the intersection of semantic web research and e-government case studies.

The symposium was organized in the following manner. The state of the art and the main challenges regarding both research areas were discussed in two keynote talks, the first by Eric Miller on the semantic web and the second by Maria Wimmer (University of Koblenz) on e-government. The state of practice in applying semantic technologies for e-government was analyzed in two invited presentations (Amith Seth, University of Georgia, and Miles Davis, TopQuadrant). The accepted presentations were organized in three main discussion areas: (1) the role of semantic technologies in publishing public services online, mainly oriented towards the usage of semantic web services for realizing public services; (2) the role of semantic technologies in building e-government applications (such as e-government portals), covering aspects of the real e-government's requirements in these applications; and (3) the specificity of the e-government domain regarding ontology engineering (for example, ontology development), focusing on the adaptations of existing methods and tools for ontologies in the e-government domain. Moreover, these topics were discussed in detail in three corresponding interest groups, which presented their findings at the end of the meeting.

The symposium showed clearly that this community exists and that additional effort is needed to make its work easier. First, better exchange of information: there are several research projects that are treating similar use cases and developing similar architectures. Second, better mutual understanding:



Reminder:

Submissions to the 2007 Spring Symposia are due October 6, 2006!

The 2007 Spring Symposium Series, will be held Monday through Wednesday, March 26–28, 2007, at Stanford University. For details consult aaai.org

more serious treatment of the e-government research in the semantic web community is needed; on the other hand, a deeper understanding of the strengths and drawbacks of semantic web technologies in the e-government community is also needed. Third, a common language is needed: locally used terminologies have to be unified (and formally represented).

The meeting ended with a very clear statement about the importance of establishing a community portal for semantic e-government, whereas the first task for the networked community would be the development of a comprehensive e-government ontology.

Andreas Abecker, Amit Sheth, Gregoris Mentzas, and Ljiljana Stojanovic served as cochairs of this symposium. The papers presented at this symposium are available as AAAI Technical Report SS-06-06.

To Boldly Go Where No Human-Robot Team Has Gone Before

Since January 2004, NASA has been pursuing a new exploration vision, the central objective of which is to establish a human-robotic presence on the Moon within the next two decades, as a precursor to exploration of Mars. Although robots have previously been used in space for scientific purposes, such as geology, future exploration robots will also have to perform a wide range of operational tasks including

assembly, construction, maintenance, and prospecting.

The purpose of this symposium was to examine novel and radical approaches to human-robot teaming, which have the potential to greatly transform space exploration. The symposium brought together researchers from a wide range of fields: robotics, mixed-initiative, human-computer interaction, dialogue systems, and intelligent agents. Three invited speakers provided diverse perspectives on human-robot interaction as applied to space exploration.

Presentations were organized around five main topics: (1) modeling (expectation management, situation monitoring, dialogue systems); (2) human-robot architecture (control schema, dynamic teams, human-aware planning); (3) user interface (demonstration learning, swarm interaction); (4) teleoperation (intent prediction, multirobots, virtual fixtures); and (5) systems (construction, prospecting, surface exploration).

A significant part of the symposium was dedicated to panel discussions. These sessions raised a number of key research questions for the human-robot interaction community to address, including: To what extent does human participation benefit space exploration? How do differences in tasks (science versus operations) influence system design, architecture, and interaction methods? What roles and responsibilities should human and robots have when working as teams? What are sensible metrics and procedures for evaluating HRI systems?

Terry Fong served as chair of this symposium. The papers presented at this symposium are available as AAAI Technical Report SS-06-07.

What Went Wrong and Why: Lessons from AI Research and Application

This symposium was dedicated to the propositions that insight often begins with unexpected results and that clarity arrives in composing a response. This perspective turns bugs, glitches, and failures into powerful instructional tools: they shape research and development by charting the boundaries of technology, revealing assumptions, and exposing design flaws.

The call for papers pursued this theme by asking authors to couple problems with discussions of lessons learned in technical, commercial, and organizational realms. The results were excellent, extracting technological lessons about feature selection, event recognition, and the performance of classification algorithms as a function of data distributions. One paper demonstrated that the real task of commercializing AI is to regularize the input data (as it required 90 percent of the code). Another described how to reconstruct false inductions (medical theories) by adding background knowledge and to infer background knowledge employed to produce unreasonable conclusions.

The symposium also attracted invited talks from seven senior figures in the field: Carl Hewitt, Craig Knoblock, Doug Lenat, John McCarthy, Mike Pazzani, Milind Tambe, and Manuela Veloso. These speakers drew an audience to match, which gave the symposium the aspect of a high-level forum on What Went Wrong and Why in AI as a whole. There was a great deal of heated and interesting discussion, but we also laughed a lot—enough to surprise the participants of adjacent symposiums.

Milind Tambe, Craig Knoblock, Manuela Veloso, and Mike Pazzani all spoke to technical themes. Tambe offered several inspiring examples about the misbehavior of the Electric Elves, including moments when the system

canceled his meeting with its sponsor and announced his location (at a coffee shop) when he had claimed to be stuck in traffic. These, and similar experiences led the ISI team to pursue a line of research in quantifying privacy loss. Knoblock described his experiences deploying a copy of the Elves at DARPA, where the tolerance for privacy violations is markedly different. Wisely, that project evolved into a related application, Travel Elves, whose goal was to ensure that a trip would execute smoothly once planned. Veloso analyzed videos of apparent mistakes by soccer bots to reveal the bones of their correct (and rather sophisticated) behavior; for example, when an Aibo head-kicked an absent ball, it became clear that it knew enough to look where the ball ought to have gone. She also argued that practical robotic problems required a mixed paradigm of formal and pragmatic methods. Mike Pazzani described his experiences commercializing AI, specifically preference-adaptive information services for the wireless web. This was a strongly cautionary tale: despite excellent technology, best product awards, and powerful customer testimonials, Adaptive-Info struggled to reach profitability and was sold in 2002. Apparently, its sin was to believe its own marketing projections, which anticipated an exponential demand curve that never materialized.

Carl Hewitt, Doug Lenat, and John McCarthy offered high-level analyses of AI. Hewitt reprised a 1985 debate, declaring that logic programs are (still) unable to encode all concurrent computations. He argued that we are all headed for a massive what went wrong and why experience in ten years, when our tools prove insufficient to program computers with 100–1000 cores on a chip. Hewitt also spoke strongly against his perceived centralization of the current funding model, with senior members of DARPA and NSF in the audience.

Lenat's talk took a more humorous tack by listing (in David Letterman style) the 12 things that have most gone wrong with AI. His list included lack of funding and loss of focus on human-level AI, but his number one item was the lack of input from com-

monsense reasoning into a wide variety of research and application tasks. This talk generated a great deal of audience participation and laughter. Ed Feigenbaum gently kept Lenat on track ("We're bored with that slide, get on to the next one!"), while McCarthy countered most every point. Lenat good-naturedly concluded with the question, "Is there anyone I haven't yet offended?" McCarthy's talk (which followed) argued forcefully that nothing has gone wrong and that AI continues to make vigorous progress. He allowed, however, that he might have been a bit ambitious in the original Dartmouth proposal.

Continuing in the belief that bugs are good, the symposium concluded with an open mike session for participants to describe their favorite what went wrong and why experiences. Among the contributions, Dan Shapiro recalled a demonstration of route planning software for the DARPA Autonomous Land Vehicle program, which proved, at great government expense, that "roads are a good idea." Carl Hewitt provided a prescriptive theory of how a field can fail to take credit for its contributions, with obvious application to AI.

In summary, this symposium provided researchers and developers with an informal, enjoyable, and extremely valuable opportunity to share their experiences about what went wrong and why. The results reaffirmed our belief that every researcher has seen bugs giving rise to insight, problems creating new research opportunities, and paths failing while teaching valuable lessons in return. These experiences should be communicated. In addition to saving the community time and effort, they capture, in some sense, the core of the research process. Said simply, if problems are the mother of necessity, they are the grandmother of invention.

Daniel Shapiro and Mehmet H. Göker served as cochairs of this symposium. The papers presented at this symposium are available as AAAI Technical Report SS-06-06.

Notes

1. At humanlevelintelligence.org.
2. caaw2006.blogspot.com.