

2003 AAAI Spring Symposium Series

Andreas Abecker, Erik K. Antonsson, Charles B. Callaway, Virginia Dignum, Patrick Doherty, Ludger van Elst, Michael Freed, Reva Freedman, Hans Guesgen, Gareth Jones, John Koza, David Kortenkamp, Mark Maybury, John McCarthy, Debasis Mitra, Jochen Renz, Debra Schreckenghost, and Mary-Anne Williams

The American Association for Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2003 Spring Symposium Series, Monday through Wednesday, 24–26 March 2003, at Stanford University.

The titles of the eight symposia were

- Agent-Mediated Knowledge Management
- Computational Synthesis: From Basic Building Blocks to High-Level Functions
- Foundations and Applications of Spatiotemporal Reasoning (FASTR)
- Human Interaction with Autonomous Systems in Complex Environments
- Intelligent Multimedia Knowledge Management
- Logical Formalization of Commonsense Reasoning
- Natural Language Generation in Spoken and Written Dialogue
- New Directions in Question-Answering Motivation

Agent-Mediated Knowledge Management

Has knowledge management research given enough attention and

The American Association for Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2003 Spring Symposium Series, Monday through Wednesday, 24–26 March 2003, at Stanford University. The titles of the eight symposia were Agent-Mediated Knowledge Management, Computational Synthesis: From Basic Building Blocks to High-Level Functions, Foundations and Applications of Spatiotemporal Reasoning (FASTR), Human Interaction with Autonomous Systems in Complex Environments, Intelligent Multimedia Knowledge Management, Logical Formalization of Commonsense Reasoning, Natural Language Generation in Spoken and Written Dialogue, and New Directions in Question-Answering Motivation.

importance to agent technology? Have agent researchers considered the potentialities and demands of the knowledge management field as an application domain? Such questions are increasingly being asked, and several projects have been started that attempt to provide an answer.

The aim of the Agent-Mediated Knowledge Management Symposium was, therefore, to bring together researchers and practitioners of both fields to discuss benefits, possibilities, and added value of cross-fertilization between knowledge management and agent technology.

Knowledge management has already been an important topic in business studies for more than a decade. From the starting days of knowledge management, technology has been recognized as an enabling, and often even a leading, factor for connecting (for example, people to other people or knowledge) and converting (such as data into knowledge). Comprehensive knowledge management endeavors, however, have always recognized that knowledge management is primarily a management science and not a computer science, implying a different role for technology in knowledge management—supporting and extending human interaction and learning—and, therefore, a need for intelligence-enhanced, integrated, and personalized solutions. Any agent researcher can tell you these goals are exactly the aims and characteristics of agents. The link is therefore established, and thus, we had the starting point for symposium participants to discuss and present their own research.

The symposium started with a keynote talk by Charles J. Petrie from Stanford University who illuminated the relationship between AI and web service technology—assets and challenges as well as drawbacks. In the symposium, 25 talks were grouped into 6 presentation sessions: (1) Collaboration and P2P Support, (2) Agent-Based Community Support, (3) Agent Models for Knowledge and Organizations, (4) Context and Personalization, (5) Ontologies and Semantic Web, and (6) Agents and Knowledge Engineering.

During these sessions, ongoing research, finished projects, and position papers from industry and academics were presented to a lively and inquisitive audience, providing nice interactions and debates both

inside and outside the conference room. Furthermore, Stanford University's outside campus facilities (together with the gorgeous California weather) provided an excellent surrounding for breakout sessions during which participants discussed comprehensive themes in small groups. Topics for these sessions were the semantic web, standardization questions, tools and methods for agent-mediated knowledge management, and further research directions. Again here, lively discussions and interaction were prevalent, and some of the groups are considering the production of articles describing their discussion and conclusions.

As the symposium evolved, it became increasingly clear that although the main premises and objectives were shared by all, different participants held different views on the field and the interaction between knowledge management and agent technology. Of course, one could a priori expect to be able to classify some participants as agent-people, who see agents as the ultimate solution for knowledge management and others as more KM-people, for which agents can be an interesting possibility for knowledge management. However, another—in AI well known—polarity soon became apparent—that between the statistical-people and the cognitive-people. The symposium therefore ended with a hilarious panel discussion for which a member of each of these four fields was asked to adopt and defend the exactly opposite view to the one he/she would usually take.

In summary, the overall consensus was that cross-fertilization between knowledge management and agent technology is a theme to be further developed because the possibilities for research and application ahead are countless. A follow-up symposium is certainly to be considered.

- Ludger van Elst
DFKI Kaiserslautern
- Virginia Dignum
University of Utrecht
- Andreas Abecker
DFKI Kaiserslautern

Computational Synthesis: From Low-Level Building Blocks to High-Level Functions

Computational synthesis research seeks generic algorithmic procedures that combine low-level building blocks in a design space to achieve a given arbitrary high-level function. For example, how do you automatically synthesize a filter circuit from basic electronic components or a locomoting robot from electromechanical components? Traditionally, there have been many knowledge-intensive approaches tailored for different problem domains. The challenge in developing more generic, domain-independent synthesis algorithms is to use knowledge-sparse methods that acquire the needed knowledge through unsupervised interaction with the problem domain. Through this interaction, synthesis algorithms will not only find a solution to the given problem but presumably will also learn something about the problem domain that will make future synthesis easier. Recently, there has been a surge of interest in these fundamental issues from three directions: (1) AI researchers interested in autonomous discovery processes; (2) engineers interested in fully automated design of increasingly complex systems; and (3) biologists interested in the origin of complexity because evolution is the primary example of a knowledge-sparse synthesis process.

The main challenge is scaling synthesis algorithms so that they can achieve complex functions, and the paths of investigation deal with automatic composition of building blocks into useful modules, automatic identification and abstraction of module function, and automatic hierarchical reuse of modules. The collection of papers presented and discussed in this symposium brought together researchers from diverse fields to exchange ideas about these common fundamental issues. The presentations fell into one of two generic approaches: (1) top-down decomposition processes that start with the high-level function and decompose it to lower levels or (2) bottom-up ap-

proaches that compose elementary building blocks into increasingly complex structures. Neither approach currently seems to successfully cross the multiscale gap spanning this exponential design space: Top down methods still require domain-dependent heuristics and sometimes user interaction, whereas bottom-up approaches cannot efficiently find compositions more than a few scales up in the hierarchy. However, several new and exciting ideas are emerging, and progress on both fronts seems likely.

The underlying assumption of presented methods is that acceptable solutions to the synthesis problem have a partially decomposable structure (although these solutions might be suboptimal). The partial subsolutions are exchangeable among different solutions and reusable across problems and can be discovered using statistical methods. The discovery of modules effectively reduces the dimensionality of the search space, albeit at the cost of limiting its versatility but possibly still retaining acceptable solutions. For example, some approaches use the high-level target criteria to identify and freeze useful compositions of building blocks, other approaches coexplore the problem space in parallel with the solution space, some approaches retain knowledge about the search space as grammars that allow more efficient exploration, and other methods take the dual approach of searching for transformations on the search space that will make the original synthesis problem easy.

- Hod Lipson
Cornell University
- Erik K. Antonsson
California Institute of Technology
- John Koza
Stanford University

Foundations and Applications of Spatiotemporal Reasoning (FASTR)

In the last few decades, tremendous progress has been made in the field of spatiotemporal knowledge management and reasoning with qualita-

tive and incomplete information. The progress has been made primarily in inventing new domains of space and time and studying complexity issues in reasoning over them. Nevertheless, there exists a lack in understanding the foundations of all these works, which is why the field has not found as much enthusiasm among the information technology practitioners as it should have. The Foundations and Applications of Spatiotemporal Reasoning (FASTR) Symposium was aimed at gaining such a fundamental understanding, focusing on three causes for the current situation: (1) fundamental, (2) methodological, and (3) strategic.

Fundamental: There is no existing generalized understanding across different domains of space and time.

Methodological: No formal general-purpose methodology has been developed across different spatiotemporal calculi studied, making it difficult to compare and contrast these disparate calculi.

Strategic: There is a lack of critical mass of application fields for each individual spatial or temporal calculi for the previous two reasons.

With a mixture of invited tutorials, short presentations of current research, and discussion sessions, the participants of the FASTR symposium spent two-and-a-half days on various issues of spatiotemporal reasoning. The start of the symposium was marked with a tutorial on qualitative spatiotemporal calculi, presented by Tony Cohn (University of Leeds), who is one of the leaders in this area. Cohn also represented the FASTR symposium at the plenary session, showing a CNN video that demonstrated in a convincing and humorous way how spatiotemporal concerns, especially the descriptions of holes, play an important role in our lives.

After the tutorial and an extensive discussion, the symposium continued with two sessions on representation and reasoning, which were filled with various short presentations and further discussions. The next day, the FASTR symposium joined the Symposium on Logical Formalization of Commonsense Reasoning in the

morning for Doug Lenat's presentation on ontological issues in the *cyc* Project, then continued with its own sessions on spatial ontologies. The afternoon of the second day was dedicated to applications and was opened by Jayant Sharma (University of Maine), who gave an invited talk on applications of spatiotemporal reasoning in location base services.

The last day of the symposium started with another highlight: Ivo Dumentsch (Brock University) presented a tutorial on relation algebras in spatial reasoning. The tutorial was followed by an introduction to the transregional collaborative research center SFB/TR8 by Thomas Barkowsky (University of Bremen). The SFB/TR8 is a major research project on spatiotemporal reasoning, which shows that major grant agencies are starting to realize the importance of the field and are willing to invest large amounts of money for research in spatiotemporal reasoning.

Although many interesting theoretical challenges were debated, a question repeatedly surfaced during the discussions: What good are all these theories about practical applications in AI and beyond? Diverging opinions emerged regarding this important issue, from the field's survivability depends on addressing the theoretical challenges to the practitioners will never need the rich formalisms developed. In a nut shell, this lively symposium was a microcosm of the prime debates within AI itself.

- Hans Guesgen
University of Auckland
- Debasis Mitra
Florida Institute of Technology
- Jochen Renz
Vienna University of Technology

Human Interaction with Autonomous Systems in Complex Environments

This symposium focused on people and automation working together to solve complex problems. The title of the symposium reflects three interacting issues: (1) humans, who are ultimately in charge of the automation;

(2) autonomous systems, which have goals and responsibilities; and (3) complex environments, especially those where mistakes have serious consequences. Bill Clancey of NASA Ames and the Institute for Human and Machine Cognition gently chided the organizers in an invited talk by suggesting the title of the symposium should more appropriately be "Agent Interaction with Human Systems in Complex Environments." That is, "start with people in their work environment and understand perceptual, cognitive, social, and environmental interactions that inhibit or facilitate work," then design autonomous agents to help.

Talks fell into broad categories that reflected the relationship of the automation with respect to humans in its environment, including autonomous systems as subordinates, where the goals and activities are defined entirely by a human supervisor; collaborators, where the goals and activities are shared with a human collaborator; tools, where the autonomous system is used by a human in a clearly defined role (for example, scheduling); and "citizens," where the autonomous system interacts with human "bystanders" as a by-product of performing its activities. Break-out groups discussed each of these categories to refine these definitions and develop metrics for evaluating how well the automation fulfills these roles.

These categories are explored briefly in the next several paragraphs.

Discussions of the subordinates' role covered applications such as autopilots for aircraft, autonomous controllers for large industrial plants and spacecraft, and surrogate software proxies representing people. Metrics for evaluating autonomous systems designed to function in this role reflected the need to not only achieve human-specified goals but also meet the distinct, practical demands of acting as an effective subordinate. For example, a good subordinate should correctly determine when it needs to ask permission before acting, when it is acceptable to interrupt its supervisor, and how to interrupt the supervisor without being disruptive.

Discussions of the autonomous systems as collaborators covered applications in which robots work alongside humans or manage information that helps humans solve complicated problems. Humans and collaborative autonomous systems have goals in common and exchange information to establish a "common ground" within which they can work together. Key research topics in this area include how the decision-making responsibility is allocated between humans and the autonomous system and how this allocation can be adjusted.

Autonomous systems as tools covered applications ranging from planning and scheduling assistants to form completion assistants. In each of these cases, the human tells the autonomous system not just what to do but also how to do it. Key features of a good autonomous tool are functions, predictability, and observability. The issue of trust and mechanisms for building trust were discussed extensively.

Autonomous systems as "citizens" included applications such as a robot wedding photographer and a robot reconnaissance soldier. In these cases, the robot encounters humans who are outside its problem focus. These people can be resources that can help the autonomous system in its task (for example, provide directions), or they might simply need to be (politely) avoided. Key topics in this area included encoding and enforcing social conventions, communicating, and dealing with perceived expectations about how the autonomous system will behave.

After working with the concept of roles throughout the symposium, it was concluded that the role that an autonomous system fills is less an indicator of inherent system capabilities and more an indicator of how the human wants to use the system. Thus, the same system can behave as a collaborator in some circumstances and as a subordinate in other circumstances. The degree of intelligence required for the various roles is not necessarily different. In many cases, the difference between an autonomous system as a tool, a subordi-

nate, and a collaborator was the spectrum of user trust in the software. The issue of trust and mechanisms for building trust are of particular importance.

The outcome of the symposium as a whole was a realization that for autonomous systems to work effectively with and alongside humans, they should be designed to do so from the start. Thus, it means more than just good interface design; it means paying attention to work practices to understand where people need help, developing trust in autonomous systems, and encoding social and organizational conventions. These will require the AI community to interact with other disciplines such as psychology, sociology, industrial engineering, human factors, and human-computer interaction.

- Michael Freed
*NASA Ames Research Center /
University of West Florida
Institute of Machine and
Human Cognition*
- David Kortenkamp
- Debra Schreckenghost
*NASA Johnson Space Center /
Metrica Inc.*

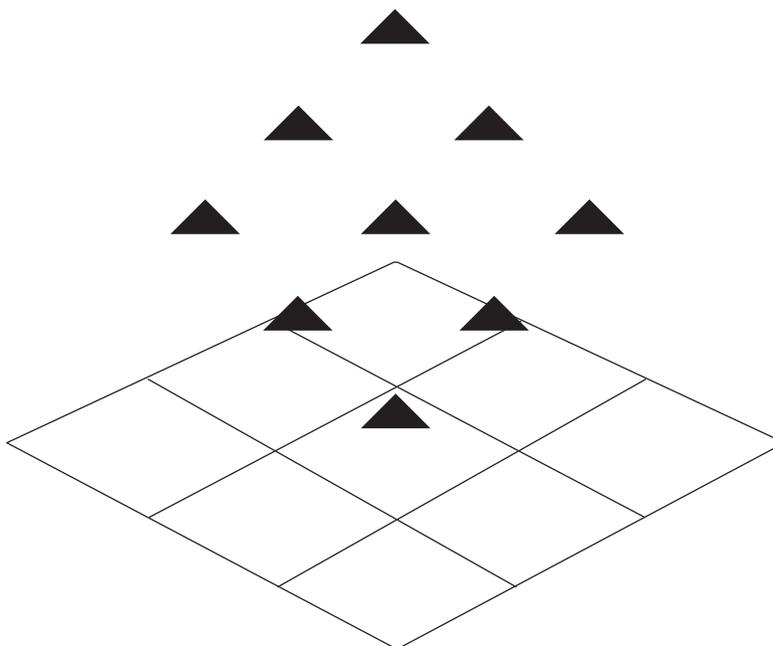
Intelligent Multimedia Knowledge Management

Research in diverse areas of media processing is starting to be integrated to form systems for intelligent management of multimedia knowledge sources. This area of work is challenging, requiring contributions from many areas of expertise. This symposium brought together researchers from multimedia systems, speech and vision technologies, visualization, information retrieval and natural language processing, and machine learning. The meeting opened with an overview of the issues in multimedia knowledge management, the technologies that can contribute to intelligent content management, and some current example systems. This overview was followed by a demonstration of the Informedia DIGITAL LIBRARY SYSTEM, developed over many years at Carnegie Mellon University (CMU), that integrates many of these techniques.

Formal presentations at the symposium were organized into themed sessions focusing on speech technologies, image processing, modeling of topics, narrative, and applications in online learning. The papers covered a broad range of work at varying levels of completion, from formative discussions of new ideas to presentations of complete working demonstration systems.

A significant factor in advancing understanding of techniques and research collaborations in various fields associated with AI has been the establishment of standard evaluation workshops. The Text-Retrieval Conference (TREC) workshop series, assessing system performance for various information-retrieval tasks, has been held annually for more than 10 years. Many of the issues explored at the symposium are found within the video-retrieval task run as part of TREC for the last three years. Alan Smeaton (Dublin City University), the coordinator of the TREC video-retrieval task, gave an invited overview of the task and its achievements. This presentation was followed by technical presentations of the work by CMU and IBM, which both participated in this evaluation task.

The symposium made time for much discussion, both in break-out sessions and between papers. Many topics were explored in these discussions, including user interaction with media systems, management of multilingual content, cross-cultural aspects of content and how this might affect knowledge management, and semantic interpretation. A strong theme that emerged in the discussions was exploring the value of video content. An important recurring question seemed to be, When does video content actually add value to multimedia content? For example, recognizing the image of a news anchor in a TV news broadcast adds nothing to our interpretation of the story. Setting this concern aside, much discussion was also devoted to the challenges of image and video processing. Much work is currently devoted to this area, with little success beyond specialized applications. We didn't solve the problems in our



**Please Join Us for the
AAAI 2004
Spring Symposium Series**

March 22 – 24, 2004

Stanford University, California

Sponsored by the
American Association for Artificial Intelligence
445 Burgess Drive, Menlo Park, CA 94025
(650) 328-3123
sss04@aaai.org

The Call for Participation is available from the AAI website
(www.aaai.org/Symposia/Spring/2004/)

lively exchanges, which produced more questions than answers, but all agreed that these were stimulating and useful discussions.

This symposium didn't provide all the answers for all the questions relating to intelligent multimedia knowledge management; indeed, it probably didn't even ask all the questions. However, attendees welcomed the open exchanges, which they stated are often not possible within more formal conferences. Many of the participants had not met previously and found this opportunity to consider the relationship between their work useful. Based on the amount of informal discussion that accompanied the meeting, it seems likely that the symposium will also achieve one of its longer-term objectives to encourage new research collaborations.

– Gareth Jones
University of Exeter

Logical Formalizations of Commonsense Reasoning

One of the major long-term goals of AI is to endow computers with commonsense reasoning capabilities. Although we know how to design and build systems that excel at certain bounded or mechanical tasks that humans find difficult, such as playing chess, we have little idea how to construct computer systems that do well at commonsense tasks that are easy for humans. Logic is a powerful modeling tool, and as a result, formalizing commonsense reasoning using logic-based approaches was the major focus; however, several papers based on other representations were presented.

The symposium was built on a long tradition of biennial meetings instigated by John McCarthy in 199X. In 2003, there were three invited speakers: (1) Marvin Minsky (Massachusetts Institute of Technology) (Minsky's talk was presented by his student Push Sing), (2) Hector Levesque (University of Toronto, Canada), and (3) Douglas Lenat (Cycorp). In addition, Leora Morgenstern (IBM) led a lively panel discussion on applications of commonsense, where contention arose over the importance

of building real systems as opposed to using rich simulations in the area of robotics and the correct way to build ontologies.

Commonsense reasoning is required in a wide variety of systems from autonomous lawn mowers to deep-space probes. Systems that exhibit commonsense have to possess robust solutions to fundamental problems in knowledge representation and reasoning. For example, an essential feature of commonsense reasoning is that it must deal with incomplete and uncertain information in dynamic environments responsively and appropriately. Furthermore, commonsense systems typically need to communicate with other systems in meaningful ways, so ontological issues naturally arise. Some new trends emerged during the symposium, for example, new attempts to integrate more traditional commonsense modeling techniques with pragmatic representations for the semantic web.

Topics of interest at the symposium included change, action, and causality; nonmonotonic reasoning and belief revision; elaboration tolerance; agents, ability, planning, and action; ontologies, including space, time, shape, matter, networks, and structures; probabilistic reasoning; belief change, update, and revision; and cognitive robotics.

– Patrick Doherty
Linköping University
– John McCarthy
Stanford University
– Mary-Anne Williams
University of Newcastle

Natural Language Generation in Spoken and Written Dialogue

It was a pleasure to participate in the symposium entitled Natural Language Generation in Spoken and Written Dialogue along with 38 fellow researchers. Participants from a variety of countries and continents described ongoing research in the intersection of natural language generation and dialogue systems from a variety of perspectives. They presented

27 papers in all— 3 long and 4 short—and hailed from the United States, United Kingdom, China, Finland, Germany, Italy, Japan, New Zealand, the Netherlands, and Sweden.

To encourage as much interaction as possible, we began the symposium with a large poster session for everyone who would later present their research articles. Thus, everyone was allowed to get to know each other as well as think of questions in advance for the presenters. Later during the conference, we also held a panel devoted to using dialogue systems in smaller university and research environments, two break-out sessions devoted to studying the deep generation of dialogue system output, and a demonstration session. Each symposium attendee helped to make this symposium a great success. Many thanks to all!

Presented papers generally fell into six categories: (1) natural language generation systems and their application to dialogue systems; (2) models of dialogue that would serve for both generation and understanding; (3) application areas with implemented systems; (4) syntactic, theoretic, and position papers; (5) descriptions of dialogue resources such as corpora and their processing; and (6) dialogue systems that had specific needs for designers of natural language-generation systems.

The papers thus described a broad range of topics, including automatic prosody markup, referring expression generation, evaluation of implemented systems, text planning and revision, conversion of existing monologue text generators to dialogue, turn taking and dialogue management, corpus-based dialogue generation, analogy generation for intelligent tutoring systems, and generation of specific linguistic phenomena such as clarification questions and topic chains. Furthermore, a number of implemented systems were described for areas such as training medical personnel with generated dialogue; creating multimodal route instructions; and driving animated, embodied agents in immersive three-dimensional worlds.

We hope the symposium has solid-



The Sixteenth Innovative Applications of Artificial Intelligence Conference

July 27-29, 2004

San Jose Convention Center, San Jose, California
Sponsored by the American Association for Artificial Intelligence

Call for Papers:

www.aaai.org/Conferences/IAAI/2004/

Paper Submission Deadline:

January 20, 2004

ified the community of researchers working at the intersection of natural language generation and dialogue by discussing shared problems and new approaches, facilitating the definition of goals and the development of shared resources, and encouraging the formation of a sense of community. We think the participants enjoyed the experience as much as we did!

- Reva Freedman
Northern Illinois University
- Charles B. Callaway
ITC-IRST

New Directions in Question Answering

Forty-seven individuals participated in the Symposium on New Directions in Question Answering. The aim of the symposium was to bring together leading scientists to report on advances in question-answering methods and technologies and together chart a road map for the future. Based on 25 submissions from the United States, Europe, Japan, and Egypt, the program committee selected 13 papers, 8, posters, and 1 panel. Three invited speakers augmented the program.

Mark Maybury (MITRE) welcomed the international group whose work covered presentations of new findings in areas including temporal question answering, multiple-perspective question answering, multimedia question answering, multilingual question answering, usability and habitability of question-answering systems, reuse in question answering, interactive- or dialogue-based question answering, question analysis, information integration, and answer presentation-generation methods and systems for question answering on the web and evaluation of question answering.

Three invited speakers presented industry, academia, and government perspectives on question answering. The first invited speaker, Brian Ulicny (Ask Jeeves) gave the keynote entitled “Putting Your Customers Questions to Work;” he described the use of dictionaries, question templates, and paraphrase authoring in

JEEVESONE 3.0, an enterprise software application for question answering. Answers might include URLs, regular expressions, and searches. Subdomain packages for domains such as finance, pharmaceuticals, and technology enhance performance. Ulicny described how JEEVESONE is context aware, for example, dealing with follow-up questions, user context, query content, or regulatory context. He contrasted JEEVESONE with TREC 10 QA tasks, pointing out their approach as site and domain specific (as opposed to open domain) focused a few key questions that matter a lot (in contrast to all questions being equal), was designed for ease of use as well as accuracy was used by nonexperts, addressed structure and unstructured data, is multilingual, has a full suite of usage analytics, and has an iterative-interactive update loop for ongoing updates and enhancements.

Deborah McGuinness, director of the Knowledge Systems Laboratory at Stanford, gave an invited talk entitled “Knowledge Representation for Question Answering.” McGuinness made the point that knowledge representation can be applied to many aspects of the question-answering problem and described both light-weight (for example, simple taxonomies, limited frame information, markup) and heavy-weight approaches to knowledge representation (for example, rich expressive languages, large hand-coded knowledge bases [for example, CYC, the Defense Advanced Research Projects Agency (DARPA) RAPID KNOWLEDGE FORMATION (RKF) and the DARPA HIGH-PERFORMANCE KNOWLEDGE BASE (HPRB)]), semiautomated generated knowledge bases, and a wide range of support tools. For example, CHIMAERA is an ontology environment tool that supports ontology analysis (for example, correctness, completeness, style), ontology term merging, and input validation. INFERENCE WEB was motivated by the need for trust and reuse in knowledge bases and provides a portable proof specification (DAML + OIL) that provides interlingua for proof interchange, a proof browser for displaying INFERENCE WEB proofs, and registry agents to record information used in proofs (for exam-

ple, sources, provenance, reasoners, rules).¹

John Prange (ARDA) gave the government-invited talk entitled “The Future of Question Answering.” Prange pointed out the need for tools to be analyst, not tool, centric. He described how his AQUAINT program focused on three key functional components: (1) question understanding and interpretation (query assessment, adviser, collaboration), (2) answer determination (for example, extracting, combining, summarizing, assessing, and inferring information), and (3) answer formulation (including refining the query). Unlike TREC QA, AQUAINT moves address the information professional, addressing a full range of questions (beyond factoid questions), contextually based question scenarios, and multiple sources and media. Prange characterized question-answering scenarios by a range of information requirements, multiple interrelated questions, source-reliability issues, and the need for background and supporting information; success is evaluated qualitatively. Finally, Prange identified some unsolved problems, including question decomposition; development of a plan for answering questions; development of reasoning and learning; knowledge of when there is no answer (but including partial, incomplete answers); and some explanations of the plan, reasoning used, and answers found. The succeeding 18-month AQUAINT workshop was held 10–12 June 2003 in San Diego, California.

A panel on web-based question answering organized by Drago Radev (University of Michigan) recognized the need to go beyond previous TREC-style question-answering tasks. This panel noted that (1) the answer of the question was known to be included in a given local corpus; (2) the size of the small corpus permitted preprocessing, including named entity extraction and the parsing of all documents; and (3) the corpus consisted of well-written news documents.

Finally, collaborative sessions within the symposium identified a road map toward a future vision of a personal assistant that would enable sce-

nario-based question answering that would support multidocument, multiperspective, and multimodal question answering. Key roadblocks to progress that require further investigation include language understanding, user need elicitation, and inference. Question answering needs cooperative, user-aware, and embedded solutions and support access to structured, semistructured, and unstructured sources. Systems can leverage a range of annotation standards (for example, TIMEML), data (for example, TIMEBANK, "Perspectives" Bank),² dictionaries (for example, WORDNET), knowledge sources (for example, CYC knowledge base), and web services (for example, JEEVES ONE, GOOGLE).

– Mark Maybury,
MITRE

Notes

1. www.ksl.stanford.edu/software/iw.
2. trec.nist.gov/data/qa.html.

Reminder:
2004 Spring
Symposium
Papers are Due
October 3, 2003.
For details please
consult the AAI
web site:
[www.aaai.org/
Symposia/Spring
/2004/](http://www.aaai.org/Symposia/Spring/2004/)

Natural Language Processing and Knowledge Representation

Language for Knowledge and Knowledge for Language



edited by
Łucja M. Iwańska
and **Stuart C. Shapiro**

Published by AAI Press
www.aaipress.org
650-328-3123



Call for Applications

Ninth AAI/SIGART Doctoral Consortium

Nineteenth National Conference on Artificial Intelligence
July 25-29, 2004 ■ San Jose, California

Sponsored by the American Association for Artificial Intelligence
and ACM/SIGART - Collocated with AAI-04

A AAI and ACM/SIGART invite students to apply for the Ninth AAI/SIGART Doctoral Consortium. The Doctoral Consortium (DC) provides an opportunity for a group of Ph.D. students to discuss and explore their research interests and career objectives with a panel of established researchers in artificial intelligence.

The consortium has the following objectives:

- To provide a setting for mutual feedback on participants' current research and guidance on future research directions
- To develop a supportive community of scholars and a spirit of collaborative research
- To support a new generation of researchers with information and advice on academic, research, industrial, and non-traditional career paths
- To contribute to the conference goals through interaction with other researchers and participation in conference events.

The Doctoral Consortium will be held as a workshop on July 25-26, 2004, immediately before the start of the main conference. Student participants in the Doctoral Consortium will receive complimentary conference registration and a fixed allowance for travel/housing.

Important Dates for Application Submission

- *February 6, 2004*: Application Package Submission Deadline
- *March 19, 2004*: Acceptance Notification
- *July 25-26, 2004*: Doctoral Consortium

The Application Packet

Applicants to the Doctoral Consortium must submit a packet consisting of six copies of the following items. Hard copy submissions are required; no electronic submissions will be accepted.

1. *Thesis Summary*. A two-page thesis summary that outlines the problem being addressed, the proposed plan for research, and a description of the progress to date. Please be sure to distinguish between work that has already been accomplished and work that remains to be done. Be sure to include a title for your work.

2. *Background Information*. Information (at most two pages) on your background and relevant experience. This should include information typically found in a curriculum vita, plus additional information that may indicate your potential contribution to the DC.

3. *Letter of Recommendation*. A letter of recommendation from your thesis advisor. It must include an assessment of the current status of your thesis research, and an *expected date for thesis submission*. In addition, your advisor should indicate what he or she hopes you would gain from participation in

the Ninth AAI/SIGART Doctoral Consortium at AAI-04.

4. *Participant's Expectations*. A short (one page or less) statement of what you expect to gain from presenting and participating in the DC, as well as what you think you can contribute to the DC.

Mail your submission packet to:

AAAI/SIGART Doctoral Consortium
445 Burgess Drive
Menlo Park, CA 94025-3442
Telephone: 650-328-3123

Review Process

The consortium organizing committee will select participants on the basis of their anticipated contribution to the workshop goals. We solicit applications from any topic area and methodology within artificial intelligence. Students will be selected who have settled on their thesis direction, but still have significant research to complete. The perfect stage is having just had a research proposal accepted by the thesis committee. Students will be selected based on clarity and completeness of the submission packet, stage of research, advisor's letter, and evidence of promise such as published papers or technical reports.

At the Conference

The organizers invite all students to attend and participate in the Doctoral Consortium, whether or not they apply to present their work. In previous years, many nonpresenting students said they found it useful to observe their peers' presentations and to participate in the ensuing discussions.

All participants selected to present their work at the Doctoral Consortium are expected to be present throughout the consortium. Our experience has been that participants gain almost as much by interacting with their peers as by having their presentations critiqued by the faculty panel. As such, we expect a commitment from participating students to attend the entire DC.

Inquiries

Additional information may be obtained by contacting the chair of the organizing committee:

Rob St. Amant
Box Campus Box 8207, 900 Main Campus Drive
Department of Computer Science
North Carolina State University
Raleigh, North Carolina 27695-8207
Telephone: (919) 515-7938
E-mail: stamant@eos.ncsu.edu