Reports on the 2005 AAAI Spring Symposium Series

Michael L. Anderson, Thomas Barkowsky, Pauline Berry, Douglas Blank, Timothy Chklovski, Pedro Domingos, Marek J. Druzdzel, Christian Freksa, John Gersh, Mary Hegarty, Tze-Yun Leong, Henry Lieberman, Ric Lowe, Susann Luperfoy, Rada Mihalcea, Lisa Meeden, David P. Miller, Tim Oates, Robert Popp, Daniel Shapiro, Nathan Schurr, Push Singh, and John Yen

■ The American Association for Artificial Intelligence presented its 2005 Spring Symposium Series on Monday through Wednesday, March 21–23, 2005 at Stanford University in Stanford, California. The topics of the eight symposia in this symposium series were (1) AI Technologies for Homeland Security; (2) Challenges to Decision Support in a Changing World; (3) Developmental Robotics; (4) Dialogical Robots: Verbal Interaction with Embodied Agents and Situated Devices; (5) Knowledge Collection from Volunteer Contributors; (6) Metacognition in Computation; (7) Persistent Assistants: Living and Working with AI; and (8) Reasoning with Mental and External Diagrams: Computational Modeling and Spatial Assistance.

The American Association for Artificial Intelligence, in cooperation with Stanford University's Computer Science Department, presented its 2005 Spring Symposium Series Monday through Wednesday, March 21–23, 2005 at Stanford University in Stanford, California. The topics of the eight symposia in this symposium series were (1) AI Technologies for Homeland Security; (2) Challenges to Decision Support in a Changing World; (3) Developmental Robotics; (4) Dialogical Robots: Verbal Interaction with Embodied Agents and Situated Devices; (5) Knowledge Collection from Volunteer Contributors; (6) Metacognition in Computation; (7) Persistent Assistants: Living and Working with AI; and (8) Reasoning with Mental and External Diagrams: Computational Modeling and Spatial Assistance.

An informal reception was held on Monday, March 21. A general plenary session, in which the highlights of each symposium were presented, was held on Tuesday, March 22. Symposia were limited to between forty and sixty participants. In the following sections, brief summaries of each symposium are presented by the symposium organizers.

AI Technologies for Homeland Security

The AI Technologies for Homeland Security symposium began with a keynote address by Robert Popp entitled "Exploiting AI, Information and Computational Social Science Technology to Understand the Adversary." Popp gave an overview of what he calls the "twenty-first century strategic threat triad," which consists of failed states, global terrorism, and proliferation of weapons of mass destruction (WMD). Popp noted that convergence of these three elements is highly destabilizing and a key strategic concern to the national security interests of the United States. He then described how various AI, information, and computational social science technologies could address various challenges associated with the strategic threat triad. In one example, he showed how information technologies, such as peer-to-peer collaboration tools, decision support structured argumentation tools, and various front-end data preprocessing tools, have helped analysts become more productive by allowing them to spend more time interpreting information and doing analysis rather than gathering information and producing reports. He also described DARPA's new preconflict anticipation and shaping initiative (PCAS), which is exploring innovative quantitative and computational social science methods and approaches—statistical, mathematical, and simulation tools—that could enable commanders and analysts to understand and anticipate the preconditions that give rise to instability and conflict within weak and failing states.

The symposium also included papers and poster presentations on several key AI technologies and their role in addressing issues regarding homeland security. Agent technologies for supporting the coordination and decision making of first responder teams were presented by researchers from Honeywell, the University of Southern California, and Pennsylvania State University. An agent-based wireless network testbed suitable for first responders was presented by Drexel University. Presentations on tools for supporting intelligence analysis included Lockheed Martin's recommender system, Northwestern's analogy-based knowledge integration, Telcordia's middleware for managing analyst teams, and Adventium's planning tool for vulnerability analysis. Techniques for analyzing terrorist networks were reported by Alphatech and the University of Arizona. The semantic web and its role for sharing and retrieving information for counter intelligence was demonstrated by Jim Hendler of the University of Maryland. The semantic web, knowledge representation tech-

niques, and natural language understanding for enhanced document retrieval were described by researchers from Stanford University, Lawrence Livermore Laboratories, SRI International, and Syracuse University. Homeland security applications for data mining and mobile robots were reported by Alphatech and the University of South Florida, respectively.

The highlights of the symposium were the two panel discussions. The panel on grand challenges touched on the challenges in leveraging works in social science and information technologies, and the challenges in deploying AI technologies (for example, the semantic web) into the hands of intelligence analysts. The second panel discussed the issue of privacy and information security. The panelist discussion included Carnegie Mellon University's privacy-aware face recognition technologies and PARC's approach for cross-source privacy appliance, which prevent queries that would allow unauthorized identification of individuals. George Cybenko (Dartmouth College) pointed out that to develop policy based on current technology while the technology is being further advanced, a continuous dialogue is critically needed between technology developers and policy makers. More questions than answers seemed to be raised by these two panels. This, perhaps, is not uncommon for an AAAI Spring Symposium.

The papers from this symposium were published in the AAAI technical report series, and are available from AAAI Press.

John Yen, Pennsylvania State University
Robert Popp, DARPA

Challenges to Decision Support in a Changing World

One of the most daunting challenges faced by decision support systems is the perpetual change in their environment. Existing decision support methodologies, tools, and frameworks are often difficult to scale up and adapt to changing knowledge, workflow, and operational setting. Adaptive systems that have to cope with change must include methodologies that go outside single theories. For example, systems that are based on probabilistic or decision-theoretic principles will be typically unable to cope with change by themselves, as neither probability theory nor decision theory says much about how the decisions are constructed, let alone how they should be modified. The general AI concepts of perception, learning, control, abstraction, and personalization must be inherently designed into the methodological, architectural, and operational aspects of adaptive systems, from application design through software and hardware infrastructure support.

The central themes for this symposium were (1) practical fielding of adaptive decision support systems, and (2) crucial technologies for successful adaptive decision support systems. The symposium included invited talks, panels, presentations, and poster sessions.

Two invited talks were given at the meeting. The first, by James F. Allen (University of Rochester), was titled "Taking Humans Seriously in Supporting Decision Making in a Complex Changing World" The second talk was given by John Doyle (North Carolina State University). Doyle's talk was called "Change and Character in Decision Support."

Panels highlighting "Challenges Posed by the Changing World" and "Bringing Problems and Methodologies Together" were also presented at the meeting.

Finally, the symposium featured presentations and poster sessions that addressed the relevant issues and technologies in the domains of biomedicine, engineering, business, military, and homeland security.

The participants were researchers with experience and interest in building adaptive decision support systems and the symposium provided a forum to highlight the various issues involved, facilitating a cross-disciplinary diffusion of methods in reasoning about and adapting to change. The presentations focused on addressing the challenges faced and the solutions involved, what worked, what did not, and why. They also aimed to chart out future research agenda by identifying specific interesting issues in various technological and application domains. The recurrent themes from the presentations included the following:

Effective user interaction. How do we build systems that can effectively support human decision makers by designing systems that can collaborate in problem solving and decision making? Discussions centered on how interfaces are crucial for user acceptance in large field decision support systems, ranging from cruise control to risk assessment in medical and manufacturing systems.

Changing preference modeling. What are the motivations and issues in representing and managing change in preferences and outcomes? Methods proposed support constructing flexible preference models, accommodating individual constraints and exhibiting "agility" for preference revision with changing environments.

Adaptive representation under uncertainty. How do we incorporate constraints and extend probabilistic representation frameworks to model changes? Approaches adopted include augmenting Bayesian networks and Markov decision processes and integrating with other formalisms to reflect decision situations that change with time, objectives, and the number and nature of knowledge sources, agents or locations.

Modeling the real world. How do we model change and develop relevant technologies motivated by real applications? Systems motivated by practical considerations often illuminate important issues and challenges that help guide the formation of new research agenda in adaptive decision support systems.

The papers from this symposium were published in the AAAI technical report series, and are available from AAAI Press.

Marek J. Druzdzel, University of Pittsburgh
Tze-Yun Leong, National University of Singapore

Developmental Robotics

The AAAI symposium on developmental robotics brought together scientists, engineers, and psychologists to discuss an approach to AI that focuses on the autonomous self-organization of general-purpose, task-independent control systems. The developmental robotics area takes its inspiration from developmental psychology and developmental neuroscience, and is a move away from task-specific methodologies where a robot is designed to solve a particular predefined goal (such as path planning to a desired location). It is different from many learning and evolutionary systems in that the reinforcement signal, teacher target, or fitness function is generated from within the system.

The symposium featured twenty presentations, seven posters, two panel presentations and much discussion. The presentations and posters included topics in philosophy, consciousness, self-awareness, anticipatory systems, value systems, reinforcement learning, and developmental architectures.

Intrinsic motivation was the topic of one of the panels, and was an issue that was discussed throughout the symposium. The central question was: What is the nature of internal pressures that will drive an autonomous system to develop increasingly sophisticated representations and behavior? There were many different approaches to answering this question, including information-theoretic methods, purely bottom-up systems, emergent motivational systems, and others.

The second panel was hosted by the qualitative reasoning research group led by Ben Kuipers from the University of Texas at Austin. They highlighted some of their work, demonstrating how a system with little innate knowledge or behavior, can bootstrap itself to more complex representations and actions. This panel presented a series of elegant and convincing proof-of-concept experiments that provide a solid foundation for the field.

Many questions emerged from the presentations. What does it mean for a stimulus to be novel, interesting, or surprising? Where in a system can a reinforcement signal come from? How can we judge if we are successful if we don't have a task with which to measure performance? Are we doing science, or is this engineering?

On the final day, we demonstrated and discussed the tools of the trade: robots, simulations, and other research software. It was decided that a repository of simulated worlds and robots would be useful in the future for comparing developmental algorithms and developed behaviors. Therefore, we have since created a community-based website[1] where one can find the beginnings of such a simulated world repository, upcoming events, and other information about Developmental Robotics.

The papers from this symposium were published in the AAAI technical report series, and are available from AAAI Press.

Doug Blank, Bryn Mawr College
Lisa Meeden (Swarthmore College)

Dialogical Robots: Verbal Interaction with Embodied Agents and Situated Devices

The aim of this symposium was to identify critical research issues in the intersection of intelligent dialog interaction and robotic behavior, and to begin developing a framework for designing and evaluating engineered systems. We chose representative application problems for focused examination, devoting a 90-minute symposium session to each. A featured presenter articulated their position statement, describing empirical data and/or engineering projects relevant to that category of human-robot dialog.

Companion robots: Human and robot engage in joint activity where the goal might be assisted grocery shopping for the visually impaired or passing time in conversation. We examined tradeoffs between hands-free spoken English dialog versus the privacy of a Braille keyboard; and the affectionate companionship of a seeing-eye dog versus inanimate dialog devices that displace human conversation.

Robots as members of human teams: We considered dialog models that combine human-robot dialog with teleoperation and human-human dialog. One suggestion treats a lunar rover as the "cerebellum" that makes local decisions to navigate terrain, collect sensor samples, and compress and transmit data to remote human operators who act as the "cortex," analyzing situations and using dialog to redirect the robot's visual system.

Teams of taskable devices: Multiple semiautonomous devices engage in collaborative task-oriented dialog. Subdialogs deal with team formation, division of labor, and team response for recovery from intermediate task failure. Architectural requirements were proposed for integrated representations and procedures to update team and individual plans, actions, dialog histories, and multiple device managers.

Robots as tour guides: Anthropomorphic robots include humanoid bipeds that deliver slide presentations and answer audience questions, talking penguins that gesture with a flapping wing, and torso robots on wheels that use arm and hand gestures to point, or facial expressions to engage individual members of a human conversation group. These dialogical robots provide artifacts for pondering and testing our intuitions about the pedagogical or affective value of the embodied tutor in a lifelike form.

Search and rescue robots: Embodied agents can exist in a shared environment as loosely or tightly coupled teams. Heterogeneous robots with diverse skills and assignments may act in concert to locate hidden objects or earthquake survivors, or to lift and vacuum under heavy furniture. This discussion raised complex issues about the relationships between perception, language interpretation, dialog management, decision-making, and physical action.

Herbert H. Clark delivered an evocative keynote address on his framework for explaining human cogni-

tive and emotional response to dialog situations involving fictional agents. If robots are a species of fictional being then, like characters in movies, plays, or video games, their emotions reside in fictional places. Dialog success rests on a tacit joint pretense that the robot has authority to perform speech acts, that is,, to welcome, warn, sympathize, and so on.

Clark's keynote set the stage for an energetic panel on robot EQ.[2] The discussion raised unresolved issues surrounding expression versus experience of emotion, the feasibility and ethics of manipulating human emotion using anthropomorphic devices, and the role of the human in projecting emotion onto inanimate object.

Susann Luperfoy, Stottler Henke Associates
David P. Miller, University of Oklahoma and KISS Institute for Practical Robotics

Knowledge Collection from Volunteer Contributors

The advent of the web has greatly simplified deployment of mass collaboration projects. Turning to volunteer contributors can allow construction of large knowledge resources. Wikipedia is a great example of a high-quality collaboratively created resource for *human* consumption; instead, we focused on creating resources aimed at *computer* consumption. The meeting addressed many aspects of this emerging direction, and provided a forum of discussions for methodologies for collecting different types on knowledge, including commonsense knowledge (ground facts, inference rules, importance and plausibility judgments); linguistic resources (paraphrase information, even associating images with word senses); and application specific knowledge, such as collection of text labels for images and correctness judgments for active learning of information extraction wrappers.

The meeting also addressed approaches to motivating users and controlling quality, including gamelike interfaces for the contribution process, rewarding corroborated contributions, and fostering contributor communities. The invited talk and the panel discussion spoke to using the collected knowledge for reasoning about everyday objects and situations, among other things. Also discussed was the potential for synergy of relying on volunteer contributions and automatic text extraction methods, with volunteers both providing seed knowledge for text extraction and reviewing, vetting and qualifying the noisy extracted knowledge.

The symposium added to the contributions of the 2003 workshop on distributed collaborative knowledge capture, and for the first time has brought together researches from many branches of AI that can both contribute to and benefit from the approach of collecting knowledge from volunteer contributors.

The symposium has attracted presentations and participation from Microsoft Research, Yahoo!, Cycorp, and researchers from universities from US, Europe and Asia.[3]

The papers from this symposium were published in the AAAI technical report series, and are available from AAAI Press.

Timothy Chklovski, University of Southern California / Information Sciences Institute
Pedro Domingos, University of Washington
Henry Lieberman, Massachusetts Institute of Technology
Rada Mihalcea, University of North Texas
Push Singh, Massachusetts Institute of Technology

Metacognition in Computation

Imagine two components, X and Y (where X and Y could be the same), related in such a way that state information flows from Y to X, and control information flows from X to Y.

Component X is in a monitoring and control relationship with component Y, and when Y is a cognitive component, we call this relationship metacognitive monitoring and control. Put formally, then, the research question we investigated in our symposium on netacognition in computation was: what are the sets $\{X, Y, S, E\}$-where Y is a cognitive component of a computational system S, and E is its environment-such that having some X in such a relationship with Y provides benefits to the system (and what are these benefits)?

Some of the hoped-for benefits were outlined in the opening keynote address on rationality and metareasoning, by Stuart Russell of the University of California, Berkeley. According to Russell, agents with limited resources having to act (and to decide how to act) in dynamic worlds need to know when they should deliberate, and when they should act. Thus, a component that could measure and monitor the value of ongoing deliberation/computation, stopping it and forcing (other) action when it ceases to have positive value, could help an agent become more rational to more effectively marshal and deploy its resources for maximal benefit. Similar benefits would be expected in the case of an agent that could monitor and control its learning components. Note how these simple examples fill in each of the variables established by our research question.

The central research question also suggests a whole host of subquestions, which were raised in the course of the symposium. For instance: how much, and what sort of state information is required for the effective monitoring of cognitive components? What are the options for the control of cognitive components—is it primarily a matter of stopping, starting and otherwise scheduling their operation, or are there effective ways to induce internal changes (for example, learning)? How much needs to be known about the inner workings of the cognitive component to effectively use or evaluate state information, and give appropriate control commands? What are the kinds of benefits we expect to see from metacognitive components, and, more importantly, how should we measure them? When and why can metacognition cause harm? Finally, do the answers to these questions depend on the details of the systems in question, or is metacognition largely domain independent?

Among the fascinating topics that emerged during the symposium were the relation between metacognition and emotion (for example, to what

degree can or should emotion be thought of as a system controlling cognition), and—because metacognition generally involves self-monitoring—the relation between metacognition and self-awareness.

To help us frame these questions and explore these topics, the symposium featured two tutorial-style keynotes in addition to the opening address, and a guided brainstorming session. John Dunlosky, a psychologist from Kent State University, outlined for us some key results in human metacognition, especially with respect to the monitoring and control of learning, and Michael Cox, of BBN Technologies, took us through the history of metacognition in computation. Tim Oates of the University of Maryland Baltimore County led the guided brainstorming session, "Whither metacognition in computation?" during which time we discussed the host of open questions in the field, and possible ways of addressing them. One small but practical upshot of this session was the establishment of a website for the field.[4]

The papers from this symposium were published in the AAAI technical report series, and are available from AAAI Press.

Michael L. Anderson, University of Maryland, College Park
Tim Oates, University of Maryland, Baltimore County

Persistent Assistants: Living and Working with AI

This symposium examined the tasks and technological barriers associated with constructing persistent assistants—intelligent agents that perform their functions while interacting with people over extended periods of time. Applications of this kind demand a richer model of agent-user interaction than has been considered to date, as they highlight the needs to establish and maintain trust, to flexibly delegate authority, to communicate state (from agent to user, and user to agent), and to tailor agent behavior to user preferences. The goal of extended operation also stresses the technology for implementing agent autonomy.

The topic of persistent assistants is interesting, in part, because there are so few examples. Microsoft's Bob™ supplied an active interface to desktop applications, but (anecdotally) it irritated users. The Electric Elves project automated tasks such as locating people and scheduling meetings, but it raised issues of privacy. During the symposium, Milind Tambe (University of Southern California) also pointed out that an early version cancelled his meeting with a research sponsor while volunteering him to give an unscheduled presentation. Illah Nourbakhsh (Carnegie Mellon University / NASA) clarified a potential reason for these difficulties in his invited talk on robotic museum guides: agents that conduct sustained interactions with people invariably become social actors and must be designed to operate on that plane. This lesson is especially relevant to mobile assistants, which invite anthropomorphic attention.

The symposium provided a good snapshot of current research on per-

sistent assistants. Adam Cheyer (SRI) described the Calo project in his invited talk—a large, and active DARPA-funded effort to construct a persistent assistant for office settings. His presentation made it clear that the end result will involve a comprehensible interface to a very complex system of competencies, knowledge, and interconnections. Fully twenty-five percent of the symposium's papers concerned Calo-motivated calendar assistants. Other applications included e-mail prioritization, briefing generation, customer complaint management, a personal disk jockey, process modeling for architecture projects, resource management for combat operations, autonomous water recovery systems, spacecraft supervision, and support for the cognitively challenged. Computational learning was a common theme, employed to tailor system behavior to user preferences. One intriguing point (raised in the disk jockey application) is that user preferences shift over time, implying that a stable policy is not necessarily the desired target of learning.

The symposium generated a number of suggestions and insights for persistent assistant design: provide rich, effective, and satisfying interactions (Illah Nourbakhsh, Carnegie Mellon University); aiding the job changes the job because users and assistants will coadapt; when learning in the wild, real-time is slow (user-assistant interactions provide sparse data for preference learning—J. and L. Gunderson, Gamma Two); and, finally "delight your users" (Jane Malin, NASA).

In summary, persistent assistants are an active area of research whose pursuit emphasizes new issues in agent design. In particular, persistence highlights the tasks of creating and maintaining the user-agent relation above and beyond the challenges of implementing agent behavior.

The papers from this symposium were published in the AAAI technical report series, and are available from AAAI Press.

Daniel Shapiro, CSLI, Stanford University
Pauline Berry, SRI International
John Gersh, Johns Hopkins University
Nathan Schurr, University of Southern California

Reasoning with Mental and External Diagrams: Computational Modeling and Spatial Assistance

This symposium brought together researchers working on the interaction between mental and external diagrams with researchers in intelligent assistance systems and in cognitive modeling for the first time. Despite their different foci these research fields carry a great potential of complementing one another to provide answers to how diagrams support intelligent reasoning processes in humans and in machines.

Mental and external diagrams form a coupled system of interacting knowledge representations. This system is effectively employed in intelligent reasoning processes, for example in spatial orientation tasks, in problems related to design and development, or in the interactive solution of complex problems. The purpose of intelligent assistance systems is to support human reasoning tasks, to enable intelligent interaction and communication with artificial systems, and to perform intelligent actions in space, for example in autonomous robot systems. Cognitive modeling aims at using software systems for reconstructing and understanding intelligent processes observed in natural cognitive systems both to provide a theoretical basis for further investigations and to develop practical applications in technical contexts.

The contributions to the symposium provided an overview over central issues currently in the focus of research. They included papers that investigate the role of animation in diagrammatic presentation of information and on human interaction with diagrammatic representations; papers that deal with problems in conducting research on spatial assistance systems or tools for assisting human spatial problem solving; papers that elaborate on the relation between external diagrams and the mental representation of graphic information; and papers that deal with general aspects of the role and function of diagrams in problem solving.

We had a highly interactive symposium with contributions by people from a wide range of disciplines: from artificial intelligence, cognitive psychology, architecture and design, cognitive robotics, geography, medicine, and from education and instruction. The symposium was thematically structured by a number of selected presentations as well as a poster session. The format of the symposium combined short plenary presentation sessions with small topical breakout sessions (in parallel) followed by plenary report-back cycles. In this way, all participants were actively involved a considerable amount of the time. The main emphasis was on producing and exchanging new ideas, perspectives, and topics for further research.

The symposium showed that the consolidation of principles and results from the three fields of research provides a rich inventory of resources which still falls short of its potential, at present. The momentum induced by this event is expected to promote research activities towards using diagrams in interactive assistance systems as well as in autonomous agents acting in spatial environments.

The papers from this symposium were published in the AAAI technical report series, and are available from AAAI Press.

Thomas Barkowsky, Universität Bremen
Christian Freksa, Universität Bremen
Mary Hegarty, University of California, Santa Barbara
Ric Lowe, Curtin University of Technology

Notes

1. www.DevelopmentalRobotics.org.

2 EQ (emotional quotient, emotional IQ, emotional intelligence, or social intelligence) refers to the ability of the robot to recognize, acknowledge, and respond appropriately to emotions, in contrast to IQ measures of cognitive intelligence.

3. For details on the presentations and links to several live knowledge collection systems, see http://teach-computers.org/kcvc05.html

4. http://www.cs.umd.edu/~anderson/MIC