AI Magazine Volume 28 Number 3 (2007) (© AAAI)

AAAI 2007 Spring Symposium Series Reports

Thomas Barkowsky, Peter Bruza, Zachary Dodds, Oren Etzioni, George Ferguson, Piotr Gmytrasiewicz, Bernhard Hommel, Benjamin Kuipers, Rob Miller, Leora Morgenstern, Simon Parsons, Holger Schultheis, Adriana Tapus, and Neil Yorke-Smith

■ The 2007 Spring Symposium Series was held Monday through Wednesday, March 26–28, 2007, at Stanford University, California. The titles of the nine symposia in this symposium series were (1) Control Mechanisms for Spatial Knowledge Processing in Cognitive/Intelligent Systems, (2) Game Theoretic and Decision Theoretic Agents, (3) Intentions in Intelligent Systems, (4) Interaction Challenges for Artificial Assistants, (5) Logical Formalizations of Commonsense Reasoning, (6) Machine Reading, (7) Multidisciplinary Collaboration for Socially Assistive Robotics, (8) Quantum Interaction, and (9) Robots and Robot Venues: Resources for AI Education.

Control Mechanisms for Spatial Knowledge Processing in Cognitive / Intelligent Systems

The purpose of this symposium was to address and investigate the interface and possible interplay between spatial knowledge processing and control processes. The former refers to the coding and use of spatial information in the perception of, the navigation in, and the communication about spatial configurations. The latter refers to all those processes that organize and integrate information, allocate processing resources, and tailor information streams to the current conditions so as to allow for coherent functioning of biological and artificial cognitive systems in their environment.

Although both areas have been researched intensely in the past, the question of how they interface has received only little attention. One reason for this could be that the two areas are modular and orthogonal to each other, so they can be readily investigated separately. And yet, essential characteristics of their interaction may be identifiable only when they are investigated together.

Either way, only a little research has explicitly addressed both this central question and the precise form of the control mechanisms involved in spatial knowledge processing. The main aim of the symposium was to bring together researchers to broach explicitly the issue of control in spatial knowledge processing for the first time.

We had a highly interactive symposium with contributions from people representing a wide range of disciplines: artificial intelligence, cognitive psychology, linguistics, neuroscience, and cognitive robotics. The symposium was thematically structured by a number of selected presentations. 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.

In the scope of the symposium, several ways of implementing control mechanisms for spatial knowledge processing were proposed, from an AI (for example, case-based reasoning), a robotics (such as reproductive perception), and a cognitive-modeling (for example, modeling the central executive) perspective. Regarding the more fundamental issue mentioned above, some participants of the symposium were advocating the view that spatial knowledge processing is controlled as any other kind of information processing, and thus, cognitive architectures such as Soar or ACT-R constitute sufficient frameworks for modeling spatial knowledge processing. Other participants, however, deemed control in spatial knowledge processing (for example, the coordination, combination, and integration of multiple spatial representations) special and not covered by existing, general-control mechanisms.

Besides these more specific results and approaches, the symposium as a whole revealed that research concerning control in spatial knowledge processing is still in its infancy. For instance, central concepts are only vaguely defined and are in need of clarification. The momentum induced by this event is expected to promote research activities toward gaining a deeper understanding of how control mechanisms for spatial knowledge processing are or should be realized in natural and artificial cognitive systems, respectively.

The papers from this symposium were published in the AAAI technical

  ISSN 0738-4602

report series and are available from AAAI Press.

—Holger Schultheis,
Thomas Barkowsky,
Benjamin Kuipers,
and Bernhard Hommel

Game-Theoretic and Decision-Theoretic Agents

This symposium marked the tenth in a series of successful game theory and decision theory symposia and workshops held over the last 10 years. The symposium attracted submissions and participation of researchers interested in principled techniques of decision and game theories to design autonomous agents. Decision theory provides a general paradigm for designing rational agents capable of operating in partially observable and nondeterministic environments. Decision-theoretic models use precise mathematical formalism to define the properties of the agent's environment, the agent's sensory capabilities, the ways the agent's actions change the state of the environment, and the agent's goals and preferences. The agent's rationality is defined as behavior that maximizes the expectation of the degree to which the preferences are achieved over time, and the planning problem is identified as a search for the optimal plan.

Game theory adds to the decision-theoretic framework the idea of multiple agents interacting within a common environment. It provides ways to specify how agents, separately or jointly, can change the environment and how the resulting changes affect their individual preferences. Building on the assumption that agents are rational and self-interested, game theory uses the notion of Nash equilibrium to design mechanisms and protocols for various forms of interaction and communication that result in the overall system behaving in a stable, efficient, and fair manner.

This year's submissions reflected the wide range of topics in planning, interacting, and learning. The discussions centered on the complementary ways the techniques of decision theory and game theory should be used to obtain designs of competent agents, learning to achieve efficient interaction or equilibria, approximate solutions to optimal planning problems, novel solution concepts, and automatic mechanism design. One of the highlights was an invited talk by Hal Varian, from the Hass Business School at the University of California at Berkeley, which illustrated the benefits of equilibrium analysis in designing auctions for position of advertisements appearing on Google.

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

—Piotr Gmytrasiewicz
and Simon Parsons

Intentions in Intelligent Systems

Intentions, in the sense of agents having specific purposes in mind when they do things, have long played a central and organizing role in the analysis of intelligent behavior. The AAAI 2007 Spring Symposium on Intentions in Intelligent Systems considered the role of intentions in implemented (or reasonably foreseeable) AI systems. The focus was primarily on practical, realistic systems that perform tasks intelligently, as opposed to abstract philosophical theories of intention or purely mathematical formalisms for representing intention. The symposium brought together key researchers from the different AI traditions to investigate intelligent systems and system behaviors based on or derived from an intentional model.

Invited speaker David Israel began the symposium with a wide-ranging and thought-provoking address on the role of belief, desire, intention (BDI) frameworks in intelligent assistants. Modeling intentional behavior using beliefs, desires, and intentions has a long history in AI (for example, the classic AAAI-84 paper by Michael Georgeff, among many others). But Israel pointed out that the "desires" component has generally not been considered in as much detail, and he noted that traditional decision theory is by and large a theory of desires in this sense. He also pointed to the relatively recent development of economic theories of "agency" (meaning an agent hired to act on behalf of a principal, for example a real estate agent) and speculated about our relationships with future intelligent agents. The discussion was at times heated, as one would hope from an opening talk.

Papers presented during the rest of the symposium covered a range of topics, from representation of and reasoning with intentions, to their role in applications as diverse as enterprise computing support, natural language dialogue, and personal assistive agents. Several papers probed the role of intentions at the level of cognitive architecture. Other papers presented investigations of the roles of intentions in human cognition, ranging from their role in infant-level word learning to the fascinating problems faced by autistic children in reasoning about the intentions of others (and what can be done to help them). The conclusion of the symposium was a joint session with the symposium on Interaction Challenges for Artificial Assistants.

So what is the role of intentions in the design and implementation of intelligent systems? A few somewhat overlapping themes emerged from the symposium, at least from this reporter's perspective. First, per David Israel's pioneering work with Michael Bratman and Martha Pollack, intentions serve to focus an agent's computational resources, given that one can't think about everything all the time. This role provides connections to decision theory and to much recent work in AI. Second, a form of "other minds" motivation: human behavior seems to involve intentions (that is, to be intentional), thus the cognitive architecture of our intelligent agents ought to involve similar notions. Third, the observation that the meaning of a wide range of commonly occurring natural language utterances seems to require both an explicit notion of intention (for content) and a process of intention recognition (for disambiguation). The SharedPlans formalism of Grosz, Sidner, and colleagues was mentioned several times in this regard. Finally, there was the interesting notion that intentions might play primarily a "design" role in the development of intelligent systems but that they might be effectively "compiled-out" of the ultimate system. This idea touches on many elements in the intellectual history of AI, from Dennett's "intentional stance" to situated automata and reactive systems. As one might therefore expect, this suggestion was not universally accepted by the symposium participants (but then again, neither were any of the others).

What did emerge from the symposium was a clear sense that intentions are already playing a fundamental role in a variety of practical, intelligent systems. Due to their central location at the intersection of a wide range of AI topics, this role is only going to increase as more complex, integrated intelligent systems are developed.

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

—George Ferguson

Interaction Challenges for Artificial Assistants

In late 2005, the comic strip OK/Cancel featured a scene titled "Clippy to the Rescue." In it, the maligned Microsoft Office Assistant, retired to the mountains of Washington, is finally vindicated. The Office Assistant as it was implemented, however, was exhibited throughout the symposium as the antithesis of successful interaction between humans and assistive agents.

This symposium was dedicated to the multifaceted challenges concerning interaction with intelligent artificial assistants. In an increasingly complex world, a new wave of such assistants has the potential to simplify and amplify our everyday personal and professional lives. These assistants will help us in mundane tasks from purchasing groceries to organizing meetings; in background tasks from providing reminders to monitoring our health; and in complex, open-ended tasks from writing a report to locating survivors in a collapsed building.

The symposium brought together practitioners and researchers of artificial intelligence, human-computer interaction, cognitive science, robotics, and assistive and agent technologies. A key thrust in fostering discussion was thematic sessions concluded by thoughts from a designated commentator, sparking subsequent discussion.

Preeminent domains were desktop assistants—participants in DARPA's Personal Assistant that Learns (PAL) program were much in evidence—and assistance for eldercare. Other applications included aerospace, tutelage, e-commerce, web services, and tourist recommendations.

Befitting a symposium concerned with interaction, participation was lively. The tone was set by the opening talk, from invited speaker Brad Myers of Carnegie Mellon University's Human Computer Interaction Institute. Myers introduced "A User Acceptance Equation for Intelligent Assistants," beginning with Clippy as an example of how to violate design principles when developing an assistive agent. With good nature, Myers admonished AI researchers for publishing statements about user acceptance or experience of systems without support from rigorous user studies; he drew the parallel of making claims about algorithm performance without empirical validation of run times.

Like Brad Myers, the second keynote speaker, Henry Lieberman of MIT's Media Laboratory, contrasted Clippy with sound practice in human-computer interaction. Lieberman presented interaction challenges for commonsense agents, also noting the opportunities that commonsense reasoning can bring when designing intelligent assistants. He urged the audience to participate in next year's Intelligent User Interfaces conference as another forum where AI and human-computer interaction HCI researchers can interact.

The lively communication of ideas continued during the poster and demonstration session and two joint sessions. The posters exemplified that effective assistance requires principled theories of collaboration, relevant AI technologies, careful interface design, and appropriate embodiment, all undergirded by user studies.

In the first joint session, with the symposium on Multidisciplinary Collaboration for Socially Assistive Robotics, a distinguished panel considered the role of physical embodiment, interaction modalities, and the degree of autonomy. Symposia participants were then challenged with an exercise in designing an eldercare robot. Panelists were kept to time by a cowbell with such effect that the bell was invited to the symposia series plenary session!

The second joint session, with the symposium on Intentions in Intelligent Systems, focused on challenges in developing agents that will be accepted in the real world. Five panelists gave their views; extended and animated discussion followed. Formal theories of collaboration and interaction and more pragmatic (or commonsense) approaches were viewed as living in complementary tension: development of intelligent agents can benefit from both the neat and the scruffy. Wayne Wobcke recounted experiences in the medical and telecommuting domains that illustrated the often surprising lessons from users.

While the focus for many contributions was in the context of software or embedded agents, intelligent assistants may be manifest as robotic embodiments. Fittingly, in addition to the first joint session, the two best papers of the symposium addressed interaction challenges in the context of embodied agents. A team from Germany considered multimodel models of interaction, while researchers from Italy reported on psychological studies for an in-home assistive robot.

Trust, learning, and modeling were significant themes that emerged. Both conceptual and application papers grappled with how to engender user trust with an agent, especially for agents that have adaptive and learning capabilities. Christopher Miller argued—against the backdrop of the PAL program—that "automated adaptiveness is not always the right answer." Participants from Japan observed that cultural difference can be significant in what users will find, to cite Myers's equation, useful, trustable, and usable. All agreed that interaction during this symposium had been a success.

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

—Neil Yorke-Smith

Logical Formalizations of Commonsense Reasoning

One of the major long-term goals of artificial intelligence is to endow computers with common sense. Although we know how to build programs that excel at certain bounded or mechanical tasks that humans find difficult, such as playing chess, we still have very little idea how to program computers to do well at commonsense tasks that are easy for humans. One approach to this problem is to formalize commonsense reasoning using mathematical logic. That was the focus of this symposium, the eighth in a series of biannual, international meetings on this topic. As with previous symposia, the symposium centered on representation rather than algorithms and on formal rather than informal methods. The symposium attracted more than 50 paper submissions, of which we were able to select just under half for presentation. The papers covered a wide range of topics in formal commonsense research, including research on commonsense domains, such as causal reasoning, planning in complex domains, reasoning about multiple agents, spatial reasoning, and physical reasoning; and research on methodological issues, such as reasoning with respect to context, reasoning with preferences, and learning and its relationship to other forms of commonsense reasoning.

The papers demonstrated how far formal commonsense reasoning has come during the last several decades. In the 1980s, researchers were still dealing with solving basic temporal reasoning problems such as the Yale Shooting Problem. Today, thanks to the completion of much foundational work, researchers can focus on more complex problems. Other advances, such as the use of simulation techniques, have helped the formal commonsense community apply their theories to real-world problems.

We were very pleased to be able to dedicate this symposium to John McCarthy, who will turn 80 in September 2007. The Commonsense'07 website includes a page of tributes to McCarthy, to which readers of this article are welcome to contribute. McCarthy, of course, is not only one of the founders of artificial intelligence, but is also the father of formal commonsense reasoning and, in fact, started this symposium series in the early 1990s. We hope this dedication went some small way to acknowledging the enormous contribution he has made to the field and the lasting leadership and vision he has provided for many of us. McCarthy attended the symposium, and we all benefited from his many insightful comments and contributions.

One of the highlights of the meeting was a panel discussion about John McCarthy's work in the area of formal commonsense reasoning. The three panelists were invited to focus on one particular seminal McCarthy paper or idea. Hector Levesque discussed "Programs with Common Sense," the 1959 paper in which McCarthy first proposed the idea of using facts written in first-order logic, along with some sort of deductive procedure, in order to automate commonsense reasoning. Levesque pointed out just how revolutionary this idea was at the time and discussed how different the approach was from anything that had been previously proposed by philosophers and logicians.

Pat Hayes spoke about "Some Philosophical Problems from the Standpoint of Artificial Intelligence," which McCarthy and Hayes coauthored in 1969. This paper, which is perhaps best known for the first detailed presentation of the situation calculus, introduced many of the foundational problems with which AI researchers are still grappling. Hayes focused on the portion of the paper discussing the distinction between various sorts of adequacy of a computational model of intelligence: metaphysical adequacy, epistemological adequacy, and heuristic adequacy. Hayes spoke about McCarthy's message, in that section of the paper, to make sure one's representation is epistemologically adequate before worrying about heuristic adequacy; and discussed existing work on the situation calculus and the frame problem relative to that theme.

Rich Thomason talked about McCarthy's 1989 paper "Artificial Intelligence, Logic, and Formalizing Common Sense." Thomason discussed the influence of McCarthy's work on the philosophical logic community and pointed out McCarthy's pioneering work in formalizing domains of knowledge and reasoning rather than, as logicians of the 19th and 20th century had done, focusing on formalizing mathematical knowledge itself.

The papers from this symposium were published in the AAAI technical report series and are available from AAAI Press as AAAI Technical Report SS-07-05.

—Rob Miller and Leora Morgenstern

Machine Reading

The time is ripe for the AI community to set its sights on machine reading—the automatic, unsupervised understanding of text. Over the last two decades or so, natural language processing has developed powerful methods for low-level syntactic and semantic text processing tasks such as parsing, semantic role labeling, and text categorization. Over the same period, the fields of machine learning and probabilistic reasoning have yielded important breakthroughs as well. It is now time to investigate how to leverage these advances to understand text.

Machine reading is very different from current semantic natural language processing research areas such as information extraction or question answering. Many natural language processing tasks utilize supervised learning techniques, which rely on hand-tagged training examples. For example, information extraction systems often utilize extraction rules learned from example extractions of each target relation. Yet machine reading is not limited to a small set of target relations. In fact, the relations encountered when reading arbitrary text are not known in advance! Thus, it is impractical to generate a set of hand-tagged examples of each relation of interest. In contrast with many natural language processing tasks, machine reading is inherently unsupervised. Another important difference is that information extraction and question answering focus on isolated "nuggets" obtained from text whereas machine reading is about forging and updating connections between beliefs. While machine reading will build on natural language processing techniques, it is a holistic process that synthesizes information gleaned from text with the machine's existing knowledge. Textual entailment is much closer in spirit to machine reading than information extraction or question answering, but with some important differences. Textual entailment systems determine whether one sentence is entailed by another. This is a valuable abstraction that naturally lends itself to tasks such as paraphrasing, summarization, and so on. Machine reading is more ambitious, however, in that it combines multiple textual entailment steps to form a coherent set of beliefs based on the text. In addition, machine reading is focused on scaling up to arbitrary relations and doing away with hand-tagged training examples. Thus, textual entailment is an important component of machine reading, but far from the whole story.

The Machine Reading symposium featured invited talks by leaders in the field including Tom Mitchell (Carnegie Mellon University), Oren Etzioni (University of Washington), and Ido Dagan (Bar Ilan University). Several themes emerged in the talks and in subsequent discussions. First, initial progress towards "reading the web" is readily apparent in systems such as TextRunner. Second, broad but shallow reading of massive corpora such as the web is very different from the focused reading process we expect when reading a scientific textbook to achieve deep understanding. Third, current approaches to the problem vary widely in the amount of background knowledge that the reader possesses at the start. Finally, inference is an integral component of the reading process and is an area of very active research as evidenced by the popularity of the textual entailment community's competitions.

We believe that this first meeting has laid the foundation for a vibrant new subfield of AI, which will continue to receive increased attention in the coming years. The field offers an avenue towards the solution of an age-old AI problem: the automatic acquisition of commonsense knowledge.

The papers from this symposium were published in the AAAI technical report series and are available from AAAI Press as AAAI Technical Report SS-07-05.

—Oren Etzioni

Multidisciplinary Collaboration for Socially Assistive Robotics

The 2007 AAAI Spring Symposium on Multidisciplinary Collaboration for Socially Assistive Robotics brought together researchers in computer science, engineering, sociology, and psychology to discuss this developing research area. Socially assistive robotics focuses on the social, rather than the physical, aspects of interaction between robots and humans in assistive contexts (such as nursing homes, hospitals, rehabilitation, and home environments). Human-robot interaction (HRI) for socially assistive applications is a growing and interdisciplinary field that draws from a range of disciplines in engineering, health sciences, social and cognitive sciences, design, and the arts. Collaboration in this domain requires close coordination and communication between diverse communities of practitioners at all stages of the process: inception, design, development, use, and evaluation. Challenges arise due to disciplinary differences in terminology, methodology, practices, and ethical considerations inherent in multidisciplinary collaboration. This symposium was a first attempt to bring these multiple and diverse research communities together in order to promote interdisciplinary discussion and knowledge transfer. Although this is a young research area, systems are currently being developed for use in hospitals, schools, and homes in therapeutic programs that monitor, encourage, and assist their users. This is an important time in the development of the field, when the broad technical community and the beneficiary populations must work together to guide the intended impact of new technologies to improve human quality of life.

The central themes for this symposium were (1) social and physical embeddedness of robots; (2) goal sharing and transfer between robots and humans; (3) nonverbal and verbal methods for establishing and maintaining the user's engagement; (4) factors relevant to the acceptance of assistive robots by a community of users (especially those with special needs); and (5) analytical frameworks and methods that can be applied to building and evaluating socially assistive robots.

The symposium featured 12 presentations, 10 posters, two panel presentations, a joint session with the Interaction Challenges for Artificial Assistants Symposium, and much discussion. The presented papers ranged from descriptions of social capabilities needed for robots to assist humans in physical or cognitive tasks (such as rehabilitation and training, therapeutic and educational play, mobility) to social factors and technical designs. The different applications for children with autism, poststroke patients, the elderly, and the visually impaired spanned a variety of technologies and robotic platforms.

The presentations identified a number of relevant issues: the role of the robot's physical embodiment; the integration of a priori knowledge about users; ensuring safety in interaction design; the relative roles played by social and physical robot assistance and how they can best be balanced; the representation of a robot's perception and competences in forms accessible and understandable to a nontechnical user; and models from psychology, cognitive science, and social science that can be utilized to advance the goals of social assistive robotics. Furthermore, in order to promote and facilitate interdisciplinary communication and discussion, symposium participants took part in two hands-on exercises where they discussed, sketched, and prototyped technologies and scenarios for assistive robots meant for different applications. We considered issues such as intent recognition, modality selection, empathy, adaptive systems, verbal and nonverbal communication, and proactive assistance.

Three invited talks were presented at the symposium. Sal Restivo (Rensselaer Polytechnic Institute, USA) discussed principles from social theory relevant to the development of socially assistive robots. Brian Scassellati (Yale University, USA) discussed the use of social robots for the diagnosis, treatment, and understanding of autism. Finally, Rachid Alami (LAAS-Toulouse, France) described work on service robots that fit in and are accepted by humans in everyday environments. Symposium participants also had the chance to visit two laboratories at Stanford University: Clifford Nass's Communication between Humans and Interactive Media laboratory, which explores the fundamental relationships between humans and interactive media; and Oussama Khatib's robotics laboratory, which focuses on safety in human-robot interaction with manipulator robotics.

Because of the novelty of this research area and high interest by participants, there are plans to continue activities in this field at future AAAI spring symposia and other venues.

The papers from this symposium were published in the AAAI technical report series and are available from AAAI Press as AAAI Technical Report SS-07-05.

—Adriana Tapus

Quantum Interaction

Just over 100 ago, quantum mechanics shocked the world of physics, explaining phenomena such as black-body radiation that had eluded classical physics. The challenging nature of quantum mechanics became part of scientific folklore, enshrined in memorable stories such as Schrödinger's Cat, Heisenberg's Uncertainty Principle, and Einstein's Wave-Particle Duality.

Seventy years ago, Garrett Birkhoff and John von Neumann discovered that parts of the logic of quantum mechanics (QM) could be expressed by considering the lattice of subspaces of a given vector space. Conjunction is expressed by standard intersection, but negation and disjunction are expressed by orthogonal complementation and linear sum operations. Not only this: the logical operators are based upon projection, introducing nonlocal and noncommutative operations.

Some 30 to 40 years ago, pioneering researchers started to build large-scale information retrieval using vector spaces, and in recent years, researchers have come to realize that the same structures used to model the logic of QM can be used to create more expressive and powerful search engines. In parallel with this development was the emergence of QM from physics into nonquantum domains such as human language, cognition, biology, organizations, and social interaction. None of this research was in the mainstream, and yet a quietly growing body of researchers was looking to QM for inspiration in addressing a spectrum of problems. In the spring of 2007, researchers came together for the first time at the AAAI symposium on "Quantum Interaction" to share results and witness a dis-

AAAI's Spring Symposium Series was held at Stanford University, Stanford California.

cipline finding its feet and bursting forth into the scientific mainstream.

The researchers attending the symposium represented contributions from a wealth of established disciplines, including logic, mathematics, psychology, economics, cultural studies, linguistics, and computer science, with artificial intelligence as the underlying theme. A common thread throughout the presentations was the use of Hilbert spaces to leverage geometric structure.

This raises many questions: how does the logic of these spaces compare to Boolean logic? What are its basic operators, what rules govern them, and what aspects of the physical world, and of human language and cognition, do they model? The first of our invited speakers, Daniel Lehmann, presented core results and motivating challenges in this area. Other mathematical results were presented, including a geometric calculus for quickly and intuitively recovering many powerful results from quantum mechanics and tensor and lattice operators that generate representations for sophisticated linguistic structures, and enable elegant streamlined solutions to old mathematical problems. As one can read in the works of the greats such as Descartes, Newton, Euler, and Gauss, an easy solution to a hard problem is an old and trusted test of success in new mathematics.

Many eye-opening applications of quantum structures were presented to the symposium. On the microscopic level, Patrick Suppes, our second invited speaker, took on the controversial topic of quantum effects in the brain. Suppes explained the progress he and his collaborators are making in designing experiments for analyzing the conditioned responses of cockroaches to single photons. This included a design to observe pairs of cockroaches reacting to individual pairs of entangled photons. Positive experimental results would provide clear evidence of the relevance of quantum phenomena to brain processes.

Directly within the macro world, presentations in the fields of linguistics, economics, and psychology demonstrated the use of quantum structures in explaining human memory, analogical reasoning, and collaboration. Of particular note, a presentation in terms of vector entanglement was used to explain experimental findings in "prisoner's dilemma" human collaboration situations in ways that were impossible given classical Bayesian assumptions. In all, the symposium's cup ran over with creative, challenging, and fascinating scientific results, theoretical and practical, some still speculative, many closely argued and clearly demonstrated, spanning many fields but with a common purpose of describing perception, cognition, and intelligence on many related levels. We have labeled this intriguing space for exploration "Quantum Interaction."

What should we conclude from these developments? Is "quantum interaction" now a fait accompli? Not yet: and while the symposium presented many algorithms that may flourish one day in quantum computers, quantum computing is no more necessary to the pursuit of quantum explanations than the space shuttle is necessary to the study of planetary orbits, and indeed, much of the mathematics we depend on was developed in the 1840s, decades before the appearance of quantum physics. In many cases, we are seeking to complement, not to replace, classical methodologies: for example, the quantum-motivated semantic space models may provide some of the necessary "flexible glue" long needed to help expert but brittle symbolic models to scale beyond their original domains of discourse. Our last invited speaker, Eleanor Rieffel, presented a challenging and fascinating lecture on the close similarities between the classical and quantum theories of probability, stressing the centrality of vector representations and the use of the tensor product to create joint or combined distributions in both theories, alerting the audience to the fact that many of the "strange" aspects of quantum mechanics are equally present in the classical theory and advising us to proceed with great caution and careful deliberation before declaring a successful experimental explanation to be necessarily the result of supposed quantum effects. In all, the symposium benefited hugely from our distinguished speakers, full of experience and knowledge and inspirational in their dedication to truth and discovery.

Where do we go from here? Are we founding a new scholarly discipline, and if so, can we even say of which branch of knowledge it should be a part? Some of the speakers had never been attracted to an artificial intelligence conference before and were surprised and delighted to discover that practitioners in AI regarded their work as deeply relevant to the modeling—and potential engineering—of intelligence and cognition. We would all like to express grateful thanks to AAAI for its support and the sheer vision to enable this interdisciplinary event to happen. Thanks also to AAAI and AFOSR for their financial support.

Where this will lead in the fullness of time, we cannot say: in the immediate coming months, we will be returning to our areas of scholarship, performing new experiments, starting new proposals and projects, and publishing research with renewed vigor, new tools, fresh insight. Fertile new collaborations hitherto unthought of are taking shape, and we will meet again next year at Oxford to discuss their fruits. At the AAAI Spring symposium of 2007, something important began.

The papers from this symposium were published in the AAAI technical report series and are available from AAAI Press as AAAI Technical Report SS-07-05.

—Peter Bruza

Robots and Robot Venues: Resources for AI Education

Curricula, platforms, venues, and outreach—these four themes wove together this symposium of AI/robotics educators and educators who use robots in their work.

The first morning highlighted participants' curriculum-building efforts, ranging from robot-themed redesigns of introductory computer science to advanced electives in which undergraduates engage in AI-related research. Break-out discussion groups then reflected on what *not* to do when designing an assignment or course to include robots.

The afternoon celebrated the remarkable innovations among robot platforms in the past year. That undergraduates are fabricating robots to compete in the RoboCup nanogram-league "2-millimeter dash" event underscores the remarkable opportunities for hands-on engineering and computational learning now available. Recent, low-cost platforms such as Lego's NXT have already found their way into undergraduate courses; their impact will only increase in the future. Indeed, the tension between building on top of this next generation of robot resources versus explicitly teaching the engineering and control challenges that underlie them was succinctly summarized in Fred Martin's paper "Real Robots Don't Drive Straight." Though the diversity of the group made full consensus impossible, appreciation for the trade-offs involved grew throughout the day's talks and full set of poster presentations and demonstrations.

The group returned Tuesday morning to consider how robot venues—both exhibitions and competitions—could be leveraged for educational ends. Several of the presentations highlighted the extent to which such venues now emphasize fundamental computational topics, for example, effective representation of tasks and sensor data, reasoning with uncertainty, and creating solutions robust to unmodeled phenomena. Through such core AI challenges, the symposium glimpsed an opportunity to create a broadly compelling "next level" of robot venues that will build on the many successful high school programs that exist today. Speakers followed their presentations by leading a full-group discussion on guiding principles and future directions for such efforts. Amid the many visions that emerged, participants united around the goal of increasing student engagement outside the walls of their home institutions by tapping into larger CS, AI, and engineering communities.

The sponsorship of the Surveyor Corporation and Road Narrows Robotics enabled the group to continue this discussion over lunch, followed by a session that looked more deeply into how educational robotics can itself help build broader educational communities. Connections spanned borders—to the Middle East and Africa—and spanned disciplines, bridging with service learning, K-12 education, art, and philosophy, to name a few. The symposium pushed its schedule to the last possible moment, with poster presenters giving demonstrations and overviews right to the start of the plenary session. David Miller then captured the symposium's essence for the plenary, punctuating it aptly with a movie clip of students taking on a robot-shaped piñata at the end of their semester's work.

The robots themselves came out on the final morning of the symposium. Participants experimented with the APIs of several new platforms: the Surveyor SRV-1, the TeRK project's Qwerkbot, KIPR's XBC controller, the Myro-based Scribbler, and iRobot's Create. Many used the "whirlwind tour" task to motivate their explorations: getting the robot first to wander as far from its initial position as possible—and then returning to the start again—all within one minute. As a finale the group gathered to cheer on (and provide obstacles for) each team's system. Scribblers chatted; Creates danced with XBCs; the Qwerkbot wandered the downstairs foyer admirably. Researchers from the University of Kent secured the top spot in the informal competition with a run of 38 robot diameters. Having ported their concurrent control language, occam-pi, to the Surveyor SRV-1 the night before, they had earned the hardware and software prizes donated by Road Narrows Robotics.

In the end, the symposium's most compelling message might have been robots' emergence as both broadly applicable and broadly accessible resources for education. The style and specifics of how physical platforms can contribute to AI, computer science, and engineering curricula vary widely, but the enthusiasm they generate—both in students and in educators—remains consistently high.

The papers from this symposium were published in the AAAI technical report series and are available from AAAI Press as AAAI Technical Report SS-07-05.

—Zachary Dodds