AAAI 2008 Fall Symposia Reports

*Jacob Beal, Paul Bello, Nicholas Cassimatis, Michael Coen,
Paul R. Cohen, Alex Davis, Mark Maybury, Alexei Samsonovich,
Andrew Shilliday, Marjorie Skubic, Joshua Taylor, Sharon Walter,
Patrick Winston, and Beverly Park Woolf*

■ *The Association for the Advancement of Artificial Intelligence was pleased to present the 2008 Fall Symposium Series, held Friday through Sunday, November 7–9, at the Westin Arlington Gateway in Arlington, Virginia. The titles of the seven symposia were (1) Adaptive Agents in Cultural Contexts, (2) AI in Eldercare: New Solutions to Old Problems, (3) Automated Scientific Discovery, (4) Biologically Inspired Cognitive Architectures, (5) Education Informatics: Steps toward the International Internet Classroom, (6) Multimedia Information Extraction, and (7) Naturally Inspired AI.*

Adaptive Agents in Cultural Contexts

The goal of the Adaptive Agents in Cultural Contexts symposium was to investigate agents and environments for behavior in social and cultural contexts and for realistic adaptation of such agents to changing situations.

Computational human behavior models, in extending a conventional information-processing approach, face two complex problems: adaptation and evolution of behavior, and the cultural specificity of cognition. These fields are vast, variegated, informed by disparate theoretical and technical disciplines, and interrelated. This symposium was intended to examine the intersection of findings from the field, theory, and applications in such areas as autonomous agent models and simulations for research, international commercial enterprise, nongovernmental organizations, and military, as well as commercial, games.

The symposium began with an invited talk given by psychologist Helen Klein of Wright State University, who addressed the dimensions of culturally specific cognition and surveyed behavioral phenomena across cultures. This theoretical basis spoke directly to one of the two major modeling approaches among the symposium's papers, focusing on behavioral aspects of the individual agent. In these approaches, cultural characteristics of behavior in individual agents was identified and implemented in the context of conventional cognitive architectures. The other major modeling approach tackled the issue of cultural change, simulating the transmission of culture through societies of agents. These approaches drew from theoretical underpinnings in genetics and memetics, studying phenomena such as coalition formation in an artificial life environment.

Applications for adaptive and culturally sensitive models included simulation-based training, decision aids, and environments for communication and collaboration between agents both automated and human controlled. Part of the symposium involved the demonstration and analysis of virtual environments for agents, along with automated agents in these environments, such as the inhabitants of an Iraqi town demonstrating culturally accurate verbal, gestural, and

psychological behavior. Other applications included case-based retrieval of narratives culturally relevant to a posed problem. Some models focused more directly on adaptation, from machine-learning and game-theoretic perspectives, but discussions suggested ways in which those adaptations might vary from one cultural context to another.

The symposium was intended not only for the presentation of research but also for bridging disciplines and connecting stakeholders among the populations of users, modelers, and researchers. To this end, there was a good deal of informal discussion and debate, particularly toward resolving disconnects—of terminology, perspective, and even of culture—between the field from which data derives, the theory developed from that data, and the applications that embody that theory. Some unexpected topics arose, such as the prospects for the use of agents and simulation environments in cultural preservation and the necessary technologies and representations to maintain them for future use. Participants expressed a willingness to continue these discussions following the weekend, to expand to other disciplines such as human-computer interaction, and to reconvene at similar such symposia in future years.

Alex Davis and Jeremy Ludwig of Stottler Henke Associates, Inc., served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-08-01.

AI in Eldercare: New Solutions to Old Problems

There are numerous challenges facing older adults as they age. AI technology offers the potential for innovative solutions, and the research community has been active in proposing solutions for these challenges. The symposium brought together researchers from eight countries representing computer science, engineering, nursing, occupational therapy, philosophy, cognitive science, cognitive neuroinformatics, and health informatics. This sympo-

sium provided an updated perspective to a previous symposium held in fall 2005 on the same topic.

One major theme of the symposium was to investigate the use of sensor networks in the home environment to provide safety, to monitor activities of daily living, to assess physical and cognitive function, and to identify changes in health that may indicate an early sign of illness. Work was also presented on using a sensor network in a nursing home setting, such as for monitoring aggression between residents. A variety of sensing modalities were discussed, including motion sensors, magnetic door sensors, accelerometers on objects, bed sensors, RFID, audio, and computer vision. Various machine-learning approaches, such as transfer learning, temporal clustering, Bayesian and Markov Model classifiers, fuzzy logic, a video codebook, and SVM classifiers, were proposed to reason about the sensor data or provide diagnosis.

Participants heard about smart home labs in which researchers and volunteers have lived for a period of time to aid in collecting sensor data. Participants also heard about an aging-in-place home in which some elders have been monitored continuously for three years, providing a wealth of sensor data. In addition, representatives from Intel and GE presented their work on sensor networks in the home, which provided an industry perspective.

The symposium also included work on smart assistive devices such as a walker that monitors gait while at the same time providing active control for mobility assistance. There were a number of robotic devices including smart wheelchairs, robotic assistants for stroke rehabilitation and music therapy, and a teddy bear robot companion. The idea of a robotic assistant in the home was explored, including the practicality and acceptance of such an assistant.

The symposium participants also discussed challenges to the research community and strategies for advancing the field. Options were explored for funding, dissemination, and the sharing of data sets. An ethics panel was included to explore some of the

ethical concerns related to the use of eldercare technology to ensure that the target population is not mistreated. The symposium ended with a brainstorming session on possible solutions for two real-life scenarios for ailing elders and their caregivers. The exercise was helpful in grounding the participants in the lives of older adults with the perspective that proposed solutions should address real needs.

The symposium was chaired by Marjorie Skubic with the help of 12 other organizers from five countries. The papers of the symposium were published as AAAI Press Technical Report FS-08-02.

Automated Scientific Discovery

Humankind has long endeavored to explain and, with the advent of AI, ultimately mechanize the overarching processes of scientific discovery. The quest begins with Aristotle's account of human deductive reasoning (the theory of the syllogism, developed to model the discoveries of Euclid) and continues through modern AI, which, through impressive systems like LT, Bacon, GT, Eurisko, and Graffiti, and many theorem provers, model finders, and computational frameworks for machine-assisted reasoning, has attained some degree of automation. Over the past 60 years, starting with AI's inaugural conference, systems such as these have automated aspects of scientific discovery. Machines have generated both novel, interesting conjectures, some of which have spawned entirely new scientific research areas, as well as the increasingly efficient techniques needed to prove or refute them.

This symposium provided an intimate setting in which principal researchers in semiautomated and fully automated scientific discovery came together, taking stock of past approaches, presenting successes and failures, and discussing new techniques and opportunities.

The three keynote addresses aligned with three larger themes in the symposium. In "Looking Both Ways," Doug Lenat (Cycorp) described his past systems AM and Eurisko, as well

as his work at Cycorp, and gave a general overview of automated scientific discovery in the past 30 years, priming discussions about lessons that ought to be learned from past systems, as well as influences that should be preserved. Another theme that emerged during the symposium emphasized the importance of ontology evolution. Jean-Gabriel Ganascia's (Paris VI University) work toward reconstructing the scientific process of Claude Bernard showed that the representation of scientific ontology as observed in historical cases can be mechanized. Alan Bundy (University of Edinburgh) and Michael Chan (University of Edinburgh) described their own work on ontology repair, demonstrating that the application of ontology repair plans can be not only formalized and mechanized but also automated. One such repair plan could, for example, resolve the zero-energy inconsistency that arises from computing the potential energy of a bouncing ball without considering the elastic energy at the time of impact.

The remainder of the symposium focused on future directions of automated scientific discovery, and attendees demonstrated various techniques under development applied to discovery. Konstantine Arkoudas (Rensselaer Polytechnic Institute) showed how genetic algorithms could discover Euler's identity, and Siemion Fajtlowicz (University of Houston) demonstrated how his Graffiti system could inspire a proof of the same theorem. Susan Epstein (Hunter College) described the psychological building blocks of scientific discovery, for example, concepts, properties, and theorems, and emphasized the importance of representation schemata capable of generating examples that satisfy certain desiderata and that are advantageous in empirical fields driven predominantly by examples. David Jensen (University of Massachusetts) described how quasi-experimental designs (methods for discovering casual relationships using existing observational data) could be mechanized and automated. Selmer Bringsjord (Rensselaer Polytechnic Institute) endeavored to show that Godel's incompleteness results, despite infamous claims to the contrary,

have not been discovered by any computing machines.

Andrew Shilliday (Rensselaer Polytechnic Institute) and Selmer Bringsjord (Rensselaer Polytechnic Institute) served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-08-03.

Biologically Inspired Cognitive Architectures

Last month, my husband and I went to a corn maze in Massachusetts. We first thought the problem should be easy and could be solved in 15 minutes. However it turned out that the maze was complicated, and it took us more than 3 hours to get through it. The maze is a good analogy to the BICA puzzle that we all are dedicated to solve, except for the fact that the BICA maze is far more complicated than the corn maze.

This quotation, from the plenary talk by Wei Chen, well characterizes the informal and exciting atmosphere of the Biologically Inspired Cognitive Architectures symposium. Trapped in a maze for three days, participants brainstormed possible ways out, trying to put together their cutting-edge models and techniques. An acceptable solution would be a route toward the goal: to re-create in a computer the most essential features of the human mind that enable its growth from a child to an adult. The challenge of creating a human-level intelligent machine was taken seriously by researchers as early as in 1955 (the Dartmouth proposal), yet today it remains a far-off dream, despite many successes in narrow tasks. The only known general solution is the human brain; therefore, the hope is that an approach based on biological inspirations will be successful.

The idea that an intelligent machine can grow up on its own like a human child is also not new in the literature (for example, science fiction), but only recently it was taken seriously, when the Defense Advanced Research Projects Agency (DARPA) launched a program called Biologically Inspired Cognitive Architectures (BICA). The goal of this program was

to capture the magic of human cognition in a computational model. It took almost a year to understand that this magic amounts to the unparalleled human ability to learn. Ironically, soon after this enlightenment, the program was terminated, but the interest in the challenge continues to grow.

The objective of this BICA symposium was to explore how the development of BICA may allow for a deeper understanding of the biological underpinnings of human cognitive and learning skills and, therefore, move the field of AI and machine learning to a breakthrough. The symposium was a mixture of paper presentations and brainstorm sessions (discussion panels). The first brainstorm objective was to identify critical ingredients that would enable cognitive growth up to an adult human level. If the analogy of a chain reaction applies to the desired bootstrapped learning process, then what is the critical mass of the cognitive chain reaction: that is, the necessary architectural components, initial intelligent capabilities, environment, and embodiment? The first panel concluded that the question, even if considered within a restricted domain of language acquisition, is more difficult than was expected a priori. In other words, we have no clear idea as to what is missing. For example, a human child may be deprived of virtually any seemingly vital ability from birth and still reach an adult level of communicational intelligence. On the other hand, embodied artifacts can see, act, talk, understand mental states, and solve specific cognitive tasks better than humans, yet for some reason do not grow up.

Discussion of the vital constraints and mechanisms that should be borrowed from biology (the second panel) was exciting and productive but also did not solve the problem. Missing components of the critical mass were unclear until the third day, when Roger Azevedo explained what is self-regulated learning and why it should be added to AI systems, and then Daniel Oblinger introduced the brain seed concept, the highlight of which is that a growing BICA needs a minimal set of fundamental drives (such as

curiosity) in order to empower the cognitive growth process. In addition, it was clarified by Art Pope that certain missing capabilities can and should be finessed at the initial stage of development.

The discussion of tests and metrics (panels 2 and 3) rejected the Turing test as a success criterion and recommended to develop and use metrics based on scalability, if the goal is a scalable cognitive growth. As a result, a clear view of a roadmap started emerging by the end of the symposium. The broader range of topics addressed by individual talks that covered other bridges of the BICA maze is available at binf.gmu.edu/~asamsono/bica.

Alexei Samsonovich served as the chair of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-08-04. Selected papers in their extended versions will be submitted for a book publication by the AAAI/MIT Press.

Education Informatics: Steps toward the International Internet Classroom

Millions of children worldwide are inadequately educated or receive no education at all. Children in the United States are educated poorly relative to children in other developed nations. Some states and demographics in the United States have achievement levels no better than in developing countries. These are not new problems and there is little hope of improvement.

Learning technologies, particularly intelligent tutoring systems, have been shown to be effective compared with the gold standard of one-on-one human tutoring. But they are expensive to build, cover narrow content, do not promote collaboration or learning in social groups, are not ubiquitous or secure, and are not available to people anytime, anywhere. Nor is it obvious that a portal for the world's intelligent tutoring systems will provide free K12 education worldwide in all the world's major languages.

While the talks in our symposium dealt with technologies, its underlying themes were more political, organizational, and focused on the perplexing question, "If we build it, will they come?" None of us doubted that tutoring systems, available content and content authoring tools, new database standards, educational data mining techniques, games technology, and huge communities of volunteers could be coordinated to build an international Internet classroom. But none of us was convinced that today's technologies are sufficiently engaging, and their knowledge sufficiently deep, to motivate the volunteers and ultimately the students. Putting it bluntly, we could not figure out how to make algebra as much fun as the World of Warcraft, Second Life, or Facebook.

The symposium had some very positive results, however. We formulated a twenty-year goal: Learning that is universally available anytime, anywhere, and free at point of use; so that average students will learn what only the best students learn today; roughly four sigma above current standards; by means of technology that enables every student to learn significantly better than when taught by master human teachers; with special emphasis on creative problem solving. The more modest two-year goal is to put all teaching resources for grades 4–12 in one place (a single portal); with easy access in multiple languages to students in multiple cultures; with 10,000 students a year visiting the site. The systems on the site should take advantage of existing content (much of it in Powerpoint), and whenever possible should be compliant with a secure standard for capturing, securing and later mining student data.

We also recognized three different approaches to content delivery: intelligent tutoring systems, game-based systems, and, for want of a better term, web browsing systems. The first two are expensive, and little is known about the efficacy of game-based systems. The latter approach—curating and guiding students intelligently to content that already exists on the web—is probably the best way to provide K–12 content, worldwide, quickly. But it cannot be our only approach because it might not be sufficiently engaging to overcome the lack of motivation in many students.

This symposium was organized by Paul Cohen, Beverly Park Woolf, and Carole Beal. Please contact any of us if you are interested in the International Internet Classroom. A followup meeting in Europe during the summer of 2009 is being planned.

Multimedia Information Extraction

Significant advances have been made in information extraction from unstructured multilingual text (see trec.nist.gov). However, increasingly large collections of audio (such as iTunes), imagery (such as Flickr), and video (such as YouTube), together with rapid and widespread growth and innovation in new information services (such as blogging, podcasting, media editing, social networking), are driving the need not only for multimedia retrieval but also for information extraction from and across heterogeneous media. The Multimedia Information Extraction symposium presented a unique opportunity to move toward an integrated view of media information extraction. The purpose of the symposium was to bring together researchers and practitioners from the language processing, speech processing, image and video processing, and spatial and temporal reasoning communities to report on current advances in multimedia information extraction theories and technologies, to foster scientific interchange among these communities and to make recommendations for future R&D direction.

The first of three invited speaker presentations introduced the Emerging Media Group (EMG) created in August 2007 by the director of National Intelligence's Open Source Center. Nancy Chinchor discussed EMG's responsibilities for specialized open source collection and exploitation of web-enabled digital media sources to include social, mobile, citizen, and virtual gaming mediums. She noted that new, unmediated communication resources are being used to influence people and that open source information can be an early indicator of something going

on in the world. As an example, local events can provide early warnings; for example, increased sales of medicine can indicate a flu outbreak.

A second invited speaker, Ruzena Bajcsy (University of California Berkeley), described human body language using environmental and body sensors. Her research recognizes and attaches labels to human activity, studies communication achieved through gestures and body language, and, eventually, aims to transfer observed body movement to robotic systems. Professor Bajcsy showed a video of dancers recorded in Illinois and California merged by teleimmersion to appear to be dancing together.

The third invited speaker, John Garofolo (NIST), noted that the gargantuan amounts of multimedia now appearing on the Internet and elsewhere are necessitating an acceleration of work in multimodal research and discussed challenges such as multimodal meeting diarization and affect and emotion recognition. Supportive discussion ensued when he discussed the value of a media wiki to foster communitywide collaboration.

A symposium subsection on language and graphic extraction included a University of Delaware, Millersville Bayesian-based high-level message recognizer for simple bar charts with text captions with 79 percent accuracy. Also demonstrated was MITRE's Audio Hot Spotting system that fuses word and telephone-based recognizers to mutually compensate for noisy meeting room data and speech and nonspeech audio analysis to detect speaker identity, indications of the speaker's emotional state, and background laughter and applause. Finally BBN presented a study of aligned named entity audio and text recognitions that shows that video text contributes a significant amount of new information beyond that which is available solely from the speech content of broadcast news videos.

The video content analysis and search section of the symposium included a Kodak research project that groups and retrieves images by location based on scene matching using recognition of unique objects and events such as birthday parties. SUNY

Buffalo research proposed using accompanying text to predict video event content.

The cross-media extraction and authoring section included a Polytechnic University presentation on an interactive cross-media prototype system that uses ontologies to enable users to explore media resources with concept-based navigation and a graphical interface for presentation integration. A DFKI presentation discussed cross-media feature extraction: using supplementary text resources to analyze and index images or video that is focused on the same topic. A Carnegie Mellon University presentation described work to automatically identify the ideological perspective of websites on social or political issues.

The poster session featured Reify Corporation, New York City College, and DFKI, and a critical challenge section included contributions from PixelForensics, AFRL, Muscle Fish, Syracuse, University of Augsburg, IRST, and SRI International to help to identify priority research gaps.

Throughout the symposium, technology strategy discussions were periodically held and directed toward evolving a research roadmap for multimodal information extraction research and defining grand challenges toward motivating community research efforts.

Many thanks go to the organizing committee and to all of the symposium participants.

Mark Maybury and Sharon Walter served as cochairs for this symposium. Papers from the symposium were published as AAAI Press Technical Report FS-08-05.

Naturally Inspired AI

The divide between how biological and computational systems solve cognitive problems and adjust to novel circumstances is readily apparent. While animals display marked flexibility in adjusting to new situations, our current computational approaches excel in well-defined, formally structured domains. This symposium brought together researchers interested in doing work that is simultaneously relevant to improving models of natural cognition

and to improving the robustness of artificial intelligent systems.

Although the participants in the Naturally Inspired AI symposium came from a broad range of backgrounds, including cognitive science, psychology, and neuroscience, the papers presented at this symposium clustered into four approximate themes: natural history, models of specific phenomena, component building, and broad models of cognition. There were many representatives of the latter group, roughly clustered into two main approaches: capability-centered models built from many cognitive specialists that reason differently (such as the PolyScheme architecture used for the work presented by Paul Bello and the architecture of competing mammalian social behaviors proposed by Ian Horswill) and knowledge-centered models that unite diverse representations through a core reasoning mechanism (such as the Companion Cognitive Systems presented by Ken Forbusm and medical patient models presented by Marjorie McShane).

Work on component building seeks to provide algorithmic tools to simplify the construction of broad models of cognition. For example, Gerald Sussman (Massachusetts Institute of Technology) presented a paper on improving the robustness of devices by forcing their inputs into familiar patterns, Jennifer Roberts (Massachusetts Institute of Technology) discussed achieving flexibility through probabalistic strategy selection, and Robert Hearn (Dartmouth College) presented a neurally inspired composable mechanism that combines supervised, unsupervised, and reinforcement learning.

Models of specific families of phenomena elucidate the computational structure of particular families of behavior without making claims to breadth. For example, Thomas Horton (North Carolina State University) presented a paper on how visual symmetries can guide tool use, John Hummel presented work on a process model of explanation, and Yaakov Gal (Massachusetts Institute of Technology) discussed how tacit knowledge can guide behavior. Broad models of cognition often attempt to demonstrate their breadth by providing equivalent im-

plementations of models of specific phenomena.

These are all fed and inspired by "natural history" investigations of the computational principles at work in natural systems. For example, Patrick Shafto (University of Louisville) presented a paper on how assuming a teacher or student relationship can accelerate learning, and Whitman Richards presented a paper on the existence of physics-based "natural modalities" connecting different aspects of perception, such as the relationship between visual and tactile roughness or scaling laws connecting body size, body shape, voice pitch, lifespan, and other factors.

Discussions at the symposium frequently returned to two themes. First, how should research on the computational nature of cognition be evaluated? Second, what are the shared precepts of researchers working in this area? Although the discussions highlighted the continuing fragmentation of research on the computational nature of cognition, there was a clear consensus that the participants shared the same goals and were interested in meeting again at similar symposia or other venues.

Jacob Beal, Paul Bello, Nicholas Cassimatis, Michael Coen, and Patrick Winston served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-08-07.

Jacob Beal is a scientist at BBN Technologies.

Paul Bello is a program manager at the Office of Naval Research.

Nicholas Cassimatis is an assistant professor in the Department of Cognitive Science at Rensselaer Polytechnic Institute.

Michael Coen is an assistant professor at School of Medicine and the Department of Computer Sciences at the University of Wisconsin–Madison.

Paul Cohen is a professor and head of the Department of Computer Science at the University of Arizona.

Alex Davis is a researcher and developer in artificial intelligence for Stottler Henke Associates, Inc.

Mark Maybury is executive director of MITRE's Information Technology Division.

Alexei Samsonovich is a research assistant professor in the Krasnow Institute for Advanced Study, George Mason University.

Andrew Shilliday is a senior software engineer at Lockheed Martin Advanced Technology Laboratories. He recently received his Ph.D. in computer science from Rensselaer Polytechnic Institute.

Marjorie Skubic is an associate professor at the University of Missouri-Columbia in the Electrical and Computer Engineering Department with a joint appointment in computer science.

Joshua Taylor is a Ph.D. student in computer science at Rensselaer Polytechnic Institute. He is affiliated with the Rensselaer Artificial Intelligence and Reasoning Laboratory, as well as the Rensselaer Tetherless World Constellation.

Sharon Walter is a computer engineer at the Air Force Research Laboratory.

Patrick Winston is Ford Professor of Artificial Intelligence and Computer Science at the Massachusetts Institute of Technology.

Beverly Park Woolf is a research professor in the Computer Science Department at the University of Massachusetts Amherst.