Reports on the AAAI 2009 Fall Symposia

Roger Azevedo, Trevor Bench-Capon, Gautam Biswas, Ted Carmichael, Nancy Green, Mirsad Hadzikadic, Oluwasanmi Koyejo, Unmesh Kurup, Simon Parsons, Roberto Pirrone, Henry Prakken, Alexei Samsonovich, Donia Scott, and Richard Souvenir

■ *The Association for the Advancement of Artificial Intelligence was pleased to present the 2009 Fall Symposium Series, held Thursday through Saturday, November 5–7, at the Westin Arlington Gateway in Arlington, Virginia. The Symposium Series was preceded on Wednesday, November 4, by a one-day AI funding seminar. The titles of the seven symposia were as follows: (1) Biologically Inspired Cognitive Architectures, (2) Cognitive and Metacognitive Educational Systems, (3) Complex Adaptive Systems and the Threshold Effect: Views from the Natural and Social Sciences, (4) Manifold Learning and Its Applications, (5) Multirepresentational Architectures for Human-Level Intelligence, (6) The Uses of Computational Argumentation, and (7) Virtual Healthcare Interaction.*

The Association for the Advancement of Artificial Intelligence was pleased to present the 2009 Fall Symposium Series, held Thursday through Saturday, November 5–7, at the Westin Arlington Gateway in Arlington, Virginia. The Symposium Series was preceded on Wednesday, November 4, by a one-day AI funding seminar. The titles of the seven symposia were as follows:

- *Biologically Inspired Cognitive Architectures*
- *Cognitive and Metacognitive Educational Systems*
- *Complex Adaptive Systems and the Threshold Effect: Views from the Natural and Social Sciences*
- *Manifold Learning and Its Applications*
- *Multirepresentational Architectures for Human-Level Intelligence*
- *The Uses of Computational Argumentation*
- *Virtual Healthcare Interaction*

An informal reception was held on Thursday, November 5. A general plenary session, in which the highlights of each symposium were presented, was held on Friday, November 6.

Biologically Inspired Cognitive Architectures

The challenge of creating a real-life computational equivalent of the human mind requires that we better understand at a computational level how natural intelligent systems develop their cognitive and learning functions. In recent years, biologically inspired cognitive architectures (BICA) have emerged as a pow-

  ISSN 0738-4602

erful new approach toward gaining this kind of understanding (here *biologically inspired* is understood broadly as "brain-mind inspired"). As a clear evidence of this trend, the symposium showcased a surprisingly broad variety of cognitive architectures and associated with them approaches to solving the challenge.

The BICA symposium was the most popular one, despite the modern field of BICA being known as a variety of disjoined communities and schools of thought that used to speak different languages and ignore each other. Naturally, almost each individual study was presented from one selected perspective. It was therefore critical for researchers representing different communities to start talking to each other, to develop a common omnidirectional view of the field, and to understand its future through comparison of alternative approaches. This objective was partially achieved in a sequence of six discussion panels, the first of which, devoted to general comparative analysis of cognitive architectures, involved 16 panelists: Drs. Lebiere (chair), Baars, Cassimatis, Chandrasekaran, Chella, Goertzel, Grossberg, Holland, Laird, Ritter, Shapiro, Stocco, Sun, Thorisson, and Wang. The panel produced a remarkable common ground: a comparative table that maps architectures onto each other and constitutes the first step toward a general comparative repository of cognitive architectures, tasks, and data. Views defended by the panelists can be classified based on the understanding of overarching goals, general methodology, and perceived difficulty, as follows.

The computational neuroscience view (represented by Stephen Grossberg) is that BICA should parsimoniously explain in detail how the brain works in the rich and unpredictable real world. To be parsimonious and efficient, the model should use a large number of simple identical elements, for example the laminar circuit design shared by all parts of the neocortex.

According to the cognitive modeling view (shared by panelists representing the ACT-R community), the goal of developing BICA is to describe computationally, and be able to predict in detail, human behavior and underlying cognitive processes. While a model of this sort is not intended to give a detailed description of the brain, it should be consistent with large-scale neuroimaging data.

According to the artificial intelligence view, the goal is to develop a practically useful artificial (general?) intelligence that will replace and outperform humans in cognitive tasks valuable to the society. When it does not matter how close to the brain-mind the solution will be, defining the critical set of capabilities becomes a hard scientific challenge. In this spirit, Ben Goertzel focused on a roadmap, Pei Wang described an approach that treats intelligence as adaptation with insufficient resources, and B. Chandrasekaran and Ron Sun called attention to cognitive functions simulated with their architectures including perceptual imagery, implicit learning, metacognition, and social psychological tasks.

According to the machine consciousness view, the goal is not intelligence per se, but generally "conscious" artifacts capable of becoming useful members of the human society. They will behave, learn, communicate, and "think" as conscious beings in general, in addition to being able to perform their specific tasks. To qualify, artifacts must inherit human values, principles of human cognition, and learning. There is no penalty for outperformance, and neurophysiological constraints do not need to be met unless they are vital for making progress. The task in this interpretation appears extremely attractive and challenging. According to Antonio Chella, we cannot hope that a single lab may be able to solve it. We need to think big and launch a large-scale project involving many U.S. and EU laboratories and funding agencies, using as models historical space projects, the linear collider at CERN, and so on. Many panelists agreed that the partners in this project should work within the same theoretical framework but not necessarily using the same architecture. Hence the need for a repository of detailed descriptions and comparative analyses of cognitive architectures.

Several additional aspects of the BICA challenge were clarified in subsequent panels: on self, on emotions, on funding (with the participation of NSF, IES, and European Commission) and on the roadmap. Discussions suggested that it is necessary to study human self and emotions in order to recreate them in artifacts; and while many funding agencies believe that with human-level intelligence we are still in the dark and can only make incremental progress within limited paradigms, the time has come to challenge this view and to propose a realistic plan for achieving the overarching goal.

Alexei Samsonovich served as the chair of the BICA symposium. The papers of the symposium were published as AAAI Press Technical Report FS-09-01.

Cognitive and Metacognitive Educational Systems

The aim of the Cognitive and Metacognitive Educational Systems AAAI symposium was to stimulate the creation of a dedicated research community about the definition of what is a metacognitive educational system. What aspects of cognition, metacognition, affect, and motivation have to be explored and integrated to achieve the goal of a new generation of metacognitive tools for enhancing learning in educational systems? Finally, what are the architectural issues to design these systems?

The idea for MCES 2009 starts from several considerations about the role of metacognition and affection in the design of a computer-based learning environment (CBLE). Intelligent tutoring systems (ITSs) are designed to support learning processes in order to facilitate the acquisition, development, use, and transfer required to solve complex tasks. Besides content management, these systems have to interact with different users and support them with several decisional processes. One of the most critical decisions includes those dealing with aspects of self-regulation. Students need to learn to regulate their learning; that is, they need to plan their learning activities, to adapt their learning strategies to meet learning goals, to become aware of changing task conditions, and to evaluate their performance. In addition, students must also regulate their affect and motivation prior to, during, and after they have used artificial systems. Teachers and other external regulating agents (for example, human tutors) have to adapt by externally regulating aspects of the environments, learning system, or the learners. Finally, human or artificial tutors have to continuously and dynamically monitor and model all of the students' activities, and to make complicated inferences about them, to ensure that learning is maximized.

The symposium hosted many contributions from researchers in heterogeneous disciplines: AI, cognitive and learning sciences, education psychology, education science, HCI, computational linguistics, web technologies, social network analysis, visualization techniques, software architectures, and multiagent systems. Discussion focused mainly on the techniques needed to stimulate and maintain a self-regulatory behavior in the learner even after the use of a CBLE. Several working systems have been presented in this respect along with some general architectural frameworks to support building such ITSs.

A stronger position has been discussed that is related to the definition of a new generation of educational systems that have to exhibit metacognition on their own to stimulate the learner. Linguistic interaction and emotions play a central role in these systems. This topic has been exploited in the final panel session, which was devoted to the relation between metacognitive educational systems and their need to become "conscious" with respect to the "machine consciousness" paradigm. The outcome of the discussion was that modern educational systems have to be designed to exhibit a sort of "conscious behavior." They are cognitive systems and have a mental state; moreover, they have to simulate emotions and to be proactive in the dialogue interaction with the user. There is no need for a true consciousness, as educational systems are not embodied like robots, so their interaction with humans has more relaxed constraints.

MCES was enriched by four keynote speeches and a joint session held with the BICA symposium. Stephen Grossberg (Department of Cognitive and Neural Systems, Boston University) discussed brain learning in health and disease as explained by the ART model. Cristina Conati (Computer Science Department and Laboratory for Computational Intelligence, University of British Columbia) gave a speech on modeling students' metacognition and affection inside an artificial tutor. Kurt VanLehn (Department of Computer Science and Engineering, Arizona State University) discussed educational systems able to stimulate the learner to retain metacognitive skills after the tutoring stops. Arthur Graesser (Department of Psychology and Institute for Intelligent Systems, University of Memphis) gave a speech on tutorial systems that are able to increase learners' metaknowledge (knowledge about cognition, emotions, and pedagogy). The joint session was devoted to future perspectives of research funding in the field of cognitive and educational systems. Some program officiers from IES, NSF, and EU were invited to illustrate the strategy of their agencies.

All the attendees appreciated the theoretical issues that emerged from the discussion of each work and from the invited talks, and they would like to attend future symposia with the same focus as this one. Roberto Pirrone, Roger Azevedo, and Gautam Biswas served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-09-02.

Complex Adaptive Systems and the Threshold Effect: Views from the Natural and Social Sciences

Threshold effects are found all around us. In economics, this could be movement from a bull market to a bear market; in sociology, it could be the spread of political dissent, culminating in rebellion; in biology, the immune response to infection or disease as the body moves from sickness to health.

Our goal was to bring together researchers from diverse fields who study these systems using the tools and techniques of complex adaptive systems (CASs). We decided to highlight threshold effects in various disciplines as one avenue toward exposing common dynamics that are found across different domains. In the past, knowledge gained in each domain about threshold effects has remained mostly exclusive to that domain, especially when the disciplines are far apart. It is our belief that by bringing together scholars who study these phenomena, we can leverage deep knowledge of one domain to gain insight into others.

Thus, our efforts here can generally be characterized as an attempt to build a community of researchers interested in explaining the same fundamental issues that cross disciplinary boundaries. This is a collaborative, multidisciplinary, and interdisciplinary endeavor that focuses on the use of CAS tools and techniques, from both the novice and expert perspectives, to shed light on these fundamental issues. For this symposium we invited three keynote speakers: Mitch Waldrop to elaborate on examples of "systems thinking," Alfred Hübler to discuss advances in high performance computing, and Bob Reynolds to present the results of a large-scale ecology simulation.

Three separate panel discussions took place, providing a productive balance between presentations and discussions among attendees. Russ Abbott chaired the panel on the fundamental connections between energy and information in CASs. Energy, in some form or another, is always needed to drive these systems. However, it is usually ignored in agent-based simulations. Is the flow of energy connected to the amount of information in a system, that is, a system's complexity? That was one of the issues discussed during this panel.

Once a model is built, it needs to be validated in some way. John Hummel convened a panel that looked both at existing policies to verify, validate, and accredit (VV&A) models and at some of the concerns particular to the complexity of CAS simulations. Richard Puddy's panel provided some real-world examples of CAS modeling, including five different efforts to use CAS models to improve child maltreatment (CM) prevention, while evaluating key factors at all levels of the social ecology: individual, family, community, and society.

Two extended workshop sessions were held in parallel the first day of the symposium. For researchers and students less familiar with CASs, Bill Rand held a tutorial on agent-based modeling. Utilizing a hands-on approach, Rand and his student volunteers taught the basics of model programming, using the NetLogo multiagent programming environment. For those already familiar with NetLogo models and simulations, Anne-Marie Grisogono and her colleagues held a workshop on Causal and Influence Networks (C&IN) in complex systems. In particular, they looked at attractors and phase spaces in causal networks, as well as methods for framing and visualizing CASs. The panel participants also presented some real-world examples, such as Beth Fulton's marine ecology model, that helped to illustrate the C&IN principals.

Spaced between these events were 18 paper presentations from many different fields that explore both CAS and threshold effects. There were presentations on politics and public policy, such as thresholds in international law (Anthony D'Amato) and simulations of community sentiments (Cathy Zanbaka). There was a session on group networks and dynamics, looking at subjects as diverse as bicycle pelotons (Hugh Trenchard), economic dynamics (Griffin Drutchas), and epistemological thresholds (Patrick Grim). Presentations from the field of biology covered everything from the immune system response to cancer (Didier Dréau) to the timing of fetus delivery (Tina Yu). Many other fields were represented as well, including sociology, linguistics, political science, and computer intelligence.

Mirsad Hadzikadic and Ted Carmichael served as the cochairs for this symposium, and the papers appear in AAAI technical report FS-09-03.

Manifold Learning and Its Applications

For many data sets (for example, images, robot sensors, biomedical data), classic analytic approaches are unsuccessful due to the high dimensionality of the data samples. To make analysis tractable, it is often necessary to find a lower-dimensional representation of the original data. One such representation is a manifold, a mathematical space that is locally euclidean. This basic definition serves as the starting point for manifold learning, which aims to model data sets as samples drawn off a manifold. The leads to a number of questions, such as: What is the best way to model a manifold? What can we do with our data after manifold learning? Is a manifold a reasonable representation of our data? Researchers from a variety of disciplines have different answers to these questions. The goal of this symposium was to identify the overlap of theory and uses of manifold learning in order to consolidate knowledge on this topic, discuss achievements in the area, and figure out the common open problems.

The Manifold Learning and Its Applications symposium was organized around four aspects of manifold learning, and each topic area contained a keynote presentation from a leading researcher, presentations of submitted papers, and an open discussion led by the keynote speaker. In Manifold Learning and Learning on Manifolds, keynote speaker Sayan Mukherjee (Duke University) discussed approaches for determining whether data has low-dimensional structure and presented work on supervised statistical modeling of manifolds. In Multimanifold Learning, Rene Vidal (Johns Hopkins University) led a discussion that focused on the cases where data is drawn from multiple (perhaps disjoint) manifolds. In Applications of Manifold Learning, Robert Pless (Washington University in St. Louis) demonstrated examples from medical image analysis where the learned manifold structure of data can be used to derive motion and deformation models. In Sparsity and Compressive Sensing, Lawrence Carin (Duke University) discussed nonparametric models for capturing low-dimensional structure in data.

There was general agreement that low-dimensional models were useful for applications such as those presented in the symposium like face recognition, robot sensor calibration, and visualization of high-dimensional data. However, there was disagreement on whether manifolds were always the best representation. Some felt that while manifolds were a convenient abstraction, it was unlikely that manifolds actually existed in many real data sets. Variants of these questions were raised throughout the symposium, and in the open discussion forums of each of the topic areas much of the conversation centered on the questions, "Do manifolds exist in real data and, if so, how do you know if your data lies on a manifold?" Consensus was not reached on some of these questions, but the participants agreed that nonlinear dimensionality reduction, manifold learning, and dimensionality estimation are all rich areas with many open questions and that future symposia, like this one, would be useful to bring together the various viewpoints from computer vision, machine learning, robotics, statistics, and all of the other disciples that develop and apply manifold learning methods.

Mikhail Belkin, Mauro Maggioni, Sridhar Mahadevan, Richard Souvenir, and Jerry Zhu served as the organizing committee of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-09-04.

Multirepresentational Architectures for Human-Level Intelligence

The goal of the Multirepresentational Architectures for Human-Level Intelligence symposium was to better understand the theory and applications of multiple representations in artificial intelligence and cognitive science.

A multiplicity of representational frameworks has been proposed for explaining and creating human-level intelligence. Each has been proven useful or effective for some class of problems but not across the board. This fact has led researchers to propose that perhaps the underlying design of cognition is multirepresentational, or hybrid, and made up of subsystems with different representations and processes interacting to produce the complexity of cognition.

Recent work in cognitive architectures has explored the design and use of such systems in high-level cognition. The main aim of this symposium was to bring together researchers who work on systems that utilize different types of representations to explore a range of questions about the theoretical framework and applications of such systems.

The symposium kicked off with a talk by B. Chandrasekaran in which he laid out some of the major distinctions in the debate surrounding multirepresentational systems. Over the course of the next couple of days, the symposium featured talks about four major cognitive architectures (Soar, ACT-R, CLARION, and Polyscheme) and their particular take on the various distinctions and debates surrounding the use of multiple representations. ACT-R, for instance, has more than one dimension along which multiple representations are realized — the symbolic versus visual or diagrammatic distinction and the symbolic versus subsymbolic distinction. Polyscheme, on the other hand, does not have a strict limit on the number of representations in the architecture. Instead, it proposes a small number of modules with each one having the opportunity to use as many representations as it needs. The only constraint is a common language that all modules must speak in order to communicate. These talks provided the fodder for discussions about the theoretical issues surrounding multirepresentational architectures.

There were also a number of talks that focused on the applications of multirepresentational systems. A talk by Susan Epstein of the City College of New York discussed the power of multiple representations for problem solving in complex domains. Her work uses a portfolio of representations that is used to represent the problem and a set of decision-making procedures that use these representations. Over the course of solving problems, the system learns which representations are better suited for the particular task. Another talk by Maithilee Kunda from Georgia Tech focused on the use of multiple representations in modeling human performance in the Raven Progressive Matrices test.

The symposium was attended by about equal numbers of theory- and application-oriented researchers and was beneficial in introducing researchers in one community to the work in the other. It was also encouraging to see the many different domains to which multirepresentational systems were being applied.

Unmesh Kurup and B. Chandrasekaran served as cochairs of this symposium. The papers of this symposium were published in the AAAI Technical Report FS-09-05.

The Uses of Computational Argumentation

Argumentation is a form of reasoning in which explicit attention is paid to the reasons for the conclusions that are drawn and how conflicts between reasons are resolved. Explicit consideration of the support for conclusions provides a mechanism, for example, to handle inconsistent and uncertain information. Argumentation has been studied both at the logical level, as a way of modeling defeasible inference, and at the dialogical level, as a form of agent interaction. Argumentation has long been studied in disciplines such as philosophy, and one can find approaches in computer science from the 1970s onward that clearly owe something to the notion of an argument. Work on computational argumentation, where arguments are explicitly constructed and compared as a means of solving problems on a computer, first started appearing in the second half of the 1980s, and argumentation is now well established as an important subfield within artificial intelligence.

We now have a good understanding of the basic requirements of argumentation systems, and there are several theoretical models that have been widely studied by researchers. We have one or two robust implementations, and the first software systems built around argumentation are beginning to appear. This, therefore, is an appropriate time to consider what these models and implementations might be used for, and this was the topic of the symposium.

The presentations and discussion at the symposium revealed that while much of the focus of argumentation research is still on refining the existing models, researchers are increasingly thinking about practical applications. Many of these applications exploit the idea that identifying the arguments—the reasons for different conclusions—helps us to better understand the problem being solved. Thus, for example, there seems to be benefit in backing automatically generated recommendations with arguments that help to explain why the recommendation was made. Similarly, in the context of group decision making, presenting the arguments can help to resolve conflicts and obtain better decisions.

However, perhaps the more interesting applications discussed were the more speculative ones. One of the strengths of argumentation is its ability to capture different points of view and give a detailed picture of how they interact. This exposes situations in which one argument seriously affects the validity of another and situations in which one argument has little to do with another. Several presenters suggested using this aspect of argumentation in the context of managing pubic debate on a topic. For example, this would provide a mechanism for structuring Internet discussions or a mechanism for government agencies to collect public opinion on policy decisions. Rather than a long list of statements, an argumentation-based discussion system would construct a structure that captured the relationships between the statements and that could then be analyzed and summarized using the formal tools that argumentation researchers have developed, allowing something close to a normative view of the state of the discussion to be formed.

Trevor Bench-Capon, Simon Parsons, and Henry Prakken served as cochairs of this event. The papers were published as AAAI technical report FS-09-06.

Virtual Healthcare Interaction

Interaction between healthcare providers and patients has a central role in patient satisfaction and successful health outcomes. Intelligent systems are beginning to play a role in this kind of interaction. They can assist in retrieval and summarization of relevant and trustworthy information and in tailoring information for the patient. Furthermore, they can provide virtual healthcare services such as reminding the patient to take his medicine, coaching a healthy lifestyle, and monitoring the patient's health. On the healthcare provider's side, they can provide virtual patients for training providers and caregivers to diagnose, care for, or develop more effective communication with clients.

The fall symposium on virtual healthcare interaction brought together health communication researchers and AI researchers and engineers from the subfields of simulation and serious games, spoken dialogue systems, natural language understanding and generation, and monitoring. Past AAAI symposia with related themes include the 2004 Fall Symposium on Dialogue Systems for Health Communication, the 2006 Spring Symposium on Argumentation for Consumers of Healthcare, and the 2008 Fall Symposium on AI in Eldercare.

The first major theme of presentations and discussion at the Virtual Healthcare interaction symposium was embodiment, that is, intelligent systems with avatars representing patients or healthcare professionals. Systems described in the presentations support user input through approaches ranging from menus and pattern recognition to full natural language understanding of typewritten or spoken input. In addition, some of the systems use AI to control other aspects of the scenario such as simulation of the patient's mental or physical state of health or external events affecting the patient. Discussion on this theme included factors contributing to fidelity and evaluating the relationship of fidelity to system effectiveness.

The second major theme was user-tailored natural language generation (NLG) of health information from the patient's medical records. An invited talk by Allen Browne and Guy Divita of the National Library of Medicine focused on the Specialist lexicon and other tools developed at NLM for natural language-processing applications that could be harnessed by NLG applications. In addition, presentations and discussion focused the considerable nontechnological “practical” challenges to NLG research in healthcare such as access to and de-identification of patient records and involvement of patients in knowledge acquisition. The third major theme was patient monitoring, that is, acquisition and use of symptom, activity, and location data, obtained automatically or by self-report. An invited talk by Marjorie Skubic highlighted monitoring technology in a senior living facility. Discussion ensued on integrating sensor data with natural language input.

At the end of the symposium the consensus of the participants was that it had been very useful to find out what others are doing in virtual healthcare interaction and that they would like to come to future symposia with a similar focus. Nancy Green and Donia Scott served as cochairs of this symposium. The papers were published as AAAI Press Technical Report FS-09-07.

Roger Azevedo is an associate professor in the Department of Psychology and Institute for Intelligent Systems, the University of Memphis.

Trevor Bench-Capon is a full professor in the Department of Computer Science at the University of Liverpool.

Gautam Biswas is a full professor in the Center for Intelligent Systems, Vanderbilt University.

Ted Carmichael is a Ph.D. candidate in the Department of Computer Science at the University of North Carolina at Charlotte.

Nancy Green is an associate professor in the Department of Computer Science at the University of North Carolina Greensboro.

Mirsad Hadzikadic is a professor in the Department of Software and Information Systems at the University of North Carolina at Charlotte and the director of the Complex Systems Institute.

Oluwasanmi Koyejo is a Ph.D. student in the Department of Electrical and Computer Engineering at the University of Texas at Austin.

Unmesh Kurup is a postdoctoral researcher in the department of cognitive science at Rensselaer Polytechnic Institute in Troy, NY.

Simon Parsons is a full professor in the Department of Computer and Information Science at Brooklyn College of the City University of New York.

Roberto Pirrone is an associate professor in the Department of Computer Science and Engineering (DINFO), University of Palermo, Italy.

Henry Prakken is an assistant professor in the computer science department at Utrecht University and a full professor of law and IT at the Law Faculty of the University of Groningen.

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

Donia Scott is a professor of computational linguistics currently a visiting professor in the Department of Informatics at the University of Sussex (UK).

Richard Souvenir is an assistant professor in the Department of Computer Science at the University of North Carolina at Charlotte.