AI Magazine Volume 25 Number 4 (2004) (© AAAI)

The 2004 AAAI Spring Symposium Series

Lola Cañamero, Zachary Dodds, Lloyd Greenwald, James Gunderson, Ayanna Howard, Eva Hudlicka, Cheryl Martin, Lynn Parker, Tim Oates, Terry Payne, Yan Qu, Craig Schlenoff, James G. Shanahan, Sheila Tejada, Jerry Weinberg, and Janyce Wiebe

■ The American Association for Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2004 Spring Symposium Series, Monday through Wednesday, March 22–24, at Stanford University. The titles of the eight symposia were (1) Accessible Hands-on Artificial Intelligence and Robotics Education; (2) Architectures for Modeling Emotion: Cross-Disciplinary Foundations; (3) Bridging the Multiagent and Multirobotic Research Gap; (4) Exploring Attitude and Affect in Text: Theories and Applications; (5) Interaction between Humans and Autonomous Systems over Extended Operation; (6) Knowledge Representation and Ontologies for Autonomous Systems; (7) Language Learning: An Interdisciplinary Perspective; and (8) Semantic Web Services. Each symposium had limited attendance. Most symposia chairs elected to create AAAI technical reports of their symposium, which are available as paperbound reports or (for AAAI members) are downloadable on the AAAI members-only Web site. This report includes summaries of the eight symposia, written by the symposia chairs.

The American Association for Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2004 Spring Symposium Series, Monday through Wednesday, March 22–24, at Stanford University. The titles of the eight symposia were (1) Accessible Hands-on Artificial Intelligence and Robotics Education; (2) Architectures for Modeling Emotion: Cross-Disciplinary Foundations; (3) Bridging the Multiagent and Multirobotic Research Gap; (4) Exploring Attitude and Affect in Text: Theories and Applications; (5) Interaction between Humans and Autonomous Systems over Extended Operation; (6) Knowledge Representation and Ontologies for Autonomous Systems; (7) Language Learning: An Interdisciplinary Perspective; and (8) Semantic Web Services. Each symposium had limited attendance. Most symposia chairs elected to create AAAI technical reports of their symposium, which are available as paperbound reports or (for AAAI members) are downloadable on the AAAI members-only Web site. This report includes summaries of the eight symposia, written by the symposia chairs.

Accessible Hands-on Artificial Intelligence and Robotics Education

This symposium grew out of the 2001 AAAI Spring Symposium on Robotics and Education. While robot platforms have played a role in artificial intelligence and robotics education for more than 30 years, the cost and size of these platforms have limited their reach. Since the previous symposium, the use of accessible, low-cost robot platforms has expanded rapidly, further promoting the adoption of robotics as a pedagogical tool. The organizers and participants were eager to share their experiences and suggestions for incorporating hands-on robotics into college-level AI courses and more general venues such as museums, science centers, and do-it-yourself Web sites.

The symposium organized contributions from more than 50 participants into sessions along six major categories: (1) curricular themes, (2) laboratory exercises, (3) hardware/software, (4) assessing approaches to AI/robotics education, (5) beyond the traditional computer science student, and (6) perspectives on AI/robotics education.

The curricular themes panel session, at which participants discussed the use of robots to teach AI concepts in small and large universities, set the stage for the diversity and depth of the symposium. Two interesting common themes emerged from this panel. First, robots provide an engaging medium for hands-on learning about intelligent agents because they are physically embodied, and second, robotics can be used to learn about different layers of intelligence beginning at the signal processing level up through the higher cognitive layers that engender immediate and long-term goals.

Presenters in the laboratory exercises sessions provided step-by-step directions for teaching AI topics using a variety of platforms. Highlights included entertaining videos covering such topics as object recognition with AIBOs and neural network control with Handy Boards, as well as live demonstrations such as dust collection with an RCX programmable, microcontroller-based brickand probabilistic localization with the ER1.

The hardware/software session considered the trade-offs in choosing platforms to support AI robotics. Presentations highlighted both commercial and custom-built systems that facilitate the use of low-cost actuation and sensing modalities in an undergraduate setting. The session concluded with a peek under the hood of the next generation of the ubiquitous Handy Board. A poster session complemented this overview by detailing participants' strategies for using such platforms in AI courses. An assessment session highlighted the efforts of participants to evaluate the impact of robotics on student learning.

The common theme of the beyond the traditional computer science student session was to highlight the mathematics and science foundations through hands-on interaction with AI robotics. These robotics adopters target audiences outside the advanced undergraduate/graduate focus of computer science and engineering departments. A wide variety of strategies included introducing robotics to nonscience majors, affecting math and science interest in high school students through mentoring, and developing processes to influence public perception of technology.

The symposium's quality and range of contributions underscored the importance of hands-on interaction with physical agents as an inspiration to future generations of AI and robotics practitioners. Follow-up activities to this symposium include a proposed special issue of this magazine.

The papers presented at this symposium are available as AAAI Technical Report SS-04-01.

– Lloyd Greenwald
Drexel University

– Zachary Dodds
Harvey Mudd College

– Ayanna Howard
Jet Propulsion Laboratory

– Sheila Tejada
University of New Orleans

– Jerry Weinberg
Southern Illinois University

Architectures for Modeling Emotion: Cross-Disciplinary Foundations

Recent years have witnessed increased interest in modeling emotion within cognitive and behavior-based agent architectures. As our understanding of the complex set of phenomena that constitute emotion grows, it is increasingly evident that continued progress necessitates focused, on-going collaborations among researchers from different disciplines, primarily AI, cognitive science, psychology, and neuroscience.

The objective of this symposium was to bring together researchers from these disciplines to discuss several core themes: (1) the need for emotion in intelligent architectures and the architectural features necessary to model emotion; (2) models of emotion elicitation via appraisal processes; (3) the effects of emotions and traits on distinct components of a cognitive architecture; (4) lessons from neuroscience; and (5) the requirements for model validation. To assure the desired level of interaction, the symposium emphasized working discussion groups, moderated panels, invited keynote addresses, and ongoing interactive poster sessions, in addition to paper presentations.

Aaron Sloman (University of Birmingham) opened the symposium with a keynote stressing the need for improved conceptual frameworks and clear definitions of emotions, closely linked to specific features and mechanisms of computational architectures. He emphasized the existence of diverse "biological minds" supporting "different classes of possible states and processes," and he underscored the need to identify the fundamental architectural building blocks that give rise to the phenomena we collectively refer to as emotions. He called for the development of a common ontology for mental architectures to improve the collaboration and comparative analysis necessary for continued progress in the field.

Two psychologists provided their perspectives on modeling appraisal and personality and integrating these within cognitive architectures. In his keynote, Craig Smith (Vanderbilt University) discussed his process model of appraisal, outlining its multiple stages, primary appraisal ("Does it matter / how?") and secondary appraisal ("What should / can I do about it?"), mediated by distinct cognitive processes (fast associative versus slower reasoning). Gerald Matthews (University of Cincinnati) discussed the effects of traits and states on cognition and behavior, highlighting the diverse cognitive mechanisms mediating these effects (such as processing and memory biases, strategy shifts), and indicating how they are distributed across architecture levels (biological, computational, and knowledge). He concluded with an assertion that "humanlike personality may require the artifact to possess a developmental history of interaction with an outside world."

The keynote by Jean-Marc Fellous (the Salk Institute for Biological Studies) presented a neuroscience perspective, highlighting the evolution of emotion theories from early sequential, localized circuitry to current complexes of parallel distributed feedback loops, with multiply-determined outcomes and no unique starting points. He suggested that the processing paradigms of emotion are fundamentally different from those underlying cognition and made several specific recommendations for implementing emotions: not localizing emotions in distinct modules, not limiting emotion generation to cognitive appraisals, and explicitly modeling temporal dynamics of emotional states. He further suggested that the neuromodulatory mechanisms that appear to "implement" emotions within the brain can be implemented in terms of architecture-wide parameter manipulations of the processes controlling agent or robot behavior.

Two moderated panels and several discussion groups addressed a variety of issues, including the necessity of emotions for intelligent architectures; limits, benefits, and approaches to modeling; and approaches and requirements for model validation.

This was the third in a series of AAAI symposia focusing on emotion, and it was gratifying to see continued progress over the past six years. The interdisciplinary audience and the focus on architecture of this year's symposium generated lively discussions, as well as a number of concrete recommendations for modeling, frameworks for analyzing emotion phenomena, and admonitions against reductionism fostered by recent developments in imaging technologies.

The papers presented at this symposium are available as AAAI Technical Report SS-04-02.

– Eva Hudlicka
Psychometrix Associates

– Lola Cañamero
University of Hertfordshire

Bridging the Multiagent and Multirobotic Research Gap

This symposium brought together researchers from the fields of multiagent systems (MAS) and multirobot systems (MRS) to identify and explore those topics that can benefit from joint research efforts. Research in MAS can be loosely characterized as the study of complex systems whose behavior can profitably be analyzed in terms of interactions between collections of elements individuated along a process-centered, or agent, dimension. Agent interactions originate from concerns over either the common good or the individual interests of agents. Within the MAS paradigm, a central area of study involves the development of algorithms through which agents can effectively coordinate their behavior. Similarly, a large segment of the multirobotic systems community is concerned with the development of robotic platforms that can work together in a coordinated fashion. At the surface, therefore, it would seem that considerable common ground should exist between the two endeavors so that much could be gained through joint research efforts.

However, research in MRS has pursued specialized methods toward solving the coordination problem. Such a state of affairs is neither healthy nor likely to be productive. The danger for MRS is that advances made in the MAS arena may unnecessarily be duplicated in the MRS world; the danger for the MAS community is that novel and important approaches put forward in MRS might be overlooked or that MRS as a domain for motivating MAS research might become marginalized.

This symposium was organized around a set of challenge problems posed by representatives of both the MRS and MAS groups, together with a collection of technical papers. At the symposium, each group presented suggested solutions to the challenge problems involving techniques such as behavior-based systems, team theories, and emergent behavior approaches. The problems and approaches were discussed in a set of panels to understand and critique alternative approaches.

– Lynne Parker
University of Tennessee

Exploring Attitude and Affect in Text: Theories and Applications

Text carries opinion, perspectives, points of view, evaluations, attitudes, emotions, and affect. To date, there has been work on classifying words as expressing positive or negative attitudes, classifying texts as factual or subjective, and classifying review articles as favorable or unfavorable.

The symposium explored deeper processing of attitudes and affect, applications, and linguistic models. In addition, it identified difficult challenges to be faced in the future.

To understand and extract attitudes from text, a natural language understanding (NLP) system must make a number of distinctions. Among those explored by participants are distinguishing attitudes from simple factual assertions, distinguishing between the author's reports from reports of other people's opinions, and distinguishing between explicitly and implicitly stated attitudes.

Applications were explored that promise to benefit from the ability to understand attitudes and texts, including indexing and retrieval of documents by opinion; automatic question answering about opinions; analysis of sentiment in the media and discussion groups toward consumer products, political issues, and so on; brand and reputation management; discovering and predicting consumer and voting trends; analyzing client discourse in therapy and counseling; determining relations between scientific texts by finding reasons for citations; generating more appropriate texts and making agents more believable; and creating writers' aids.

Difficult challenges remain, however. Participants argued that analyzing attitude and affect in text is an "NLP"-complete problem. The interpretation of attitude and affect depends on audience, context, and world knowledge. In addition, there is much yet to learn about the psychological and biological relationships between emotion and language.

To continue to progress in this area in NLP, more comprehensive theories of emotion, attitude, and opinion are needed, as are lexicons of affective terms, knowledge of how such terms are used in context, and annotated corpora for training and evaluation.

This symposium, a first foray into this area, judging by the submission level (39 submissions and three invited panel discussions), number of registered participants (more than 60), and general feedback, was a huge success.

The papers presented at this symposium are available as AAAI Technical Report SS-04-07.

– Yan Qu
Clairvoyance Corporation

– James G. Shanahan
Clairvoyance Corporation

– Janyce Wiebe
University of Pittsburgh

Interaction between Humans and Autonomous Systems over Extended Operation

In 2003, the Symposium on Human Interaction with Autonomous Systems in Complex Environments concluded that if autonomous systems are going to work effectively with and alongside humans, they should be designed to do so from the start. This year's symposium extended this topic, recognizing that many future applications for autonomous systems require long-term deployment in order to achieve the desired payoff. Long-term deployment both increases the total amount of interaction between humans and autonomous systems and changes the nature of these interactions. Over extended operation, equipment will degrade, human preferences and needs will change, environmental context will drift, novel influences and obstacles will appear for which an autonomous system has not been explicitly prepared, and high-level goals and missions may shift. To be effective over extended operation, autonomous systems must be designed for appropriate interaction with humans to realize and respond to these changes.

Our group was equally divided among government/industry and academics and came from a diverse set of backgrounds including computer science, engineering, cognitive science, philosophy, psychology, human factors, and physics. From the very beginning, our symposium was characterized by a great deal of interaction and spirited debate among the participants. The opening talk by David Woods from Ohio State University drew out several issues including the practical limits of adaptation and the relative fragility of current autonomous systems (for example, robots), which frustrates progress in studying how they interact with humans. His presentation engendered a great debate about the appropriate and necessary role for humans in interactions with autonomous flight management systems. On one side, because the world is unpredictable and designers cannot possibly project all the things an autonomous system might need to know, model, and respond to, the human was assigned the role of bridging the context gap and guiding adaptation. On the other side, the domain of flight management (such as autopilot) was thought to be predictable enough that humans should let the automation do its job (at which it performs better than humans) without interfering. Participants on both sides of the debate related several crash-and-burn stories to support their point of view and were still arguing through the bitter end of the symposium.

In addition to flight management, we considered a wide range of application areas including guides for the visually impaired, reminding systems for elder care, control of advanced life-support systems, tracking human activities, personal calendar assistants, monitoring space satellite science goals, oil refinery monitoring and control, autonomous tour guide systems, unmanned autonomous vehicles, and support for anomaly response teams. We also considered issues of theory such as multiagent learning, following route directions, short-term versus long-term memory implementation, knowledge representation, coordinating teams of humans and robots, task-oriented dialogue, event pattern recognition, image retrieval, and even quantum models of autonomy.

This diverse group of interests provided a solid foundation for our panel discussion exploring the role of adaptation, "Is the ability to adapt a necessary condition for successful long-term interaction between humans and autonomous systems?" Those who said yes believed that unless designers are omniscient it is not possible to maintain a successful long-term relationship between a human, who adapts and changes, and an autonomous system, which does not adapt. Those who said no favored predictability and transparency over the ability to adapt in order to ensure enduring success. Some said, "it depends," indicating that noncomplex systems may not need to adapt but that there may be a threshold of complexity above which adaptation becomes necessary for effective operation. Many people agreed that if a system adapts, it should do so in a way that the human expects and understands to create successful interaction.

Our final panel explored the state of the art for human interaction with autonomous systems that are deployed over long durations. Of the research presented at the symposium, over half represented implemented systems, about one-third represented systems that had interacted with at least one user who was not a system developer, and ~15 percent represented systems interacting with at least one user for at least one month. We determined that the primary barriers to progress in this emerging field are (1) the large investments of time and money required to bring a deployed system to life and keep it running while we work to study how it should interact with humans and (2) the tension between the current demand for applications and real-world experience with these systems and the need to build and study autonomous systems in a safe environment where failure does not have catastrophic consequences.

The papers presented at this symposium are available as AAAI Technical Report SS-04-03.

– James Gunderson
Gamma Two, Inc.

– Cheryl Martin
NASA Johnson Space Center

Knowledge Representation and Ontologies for Autonomous Systems

This symposium was motivated by the desire to bring together experts in the autonomous systems, knowledge representation, ontology, and data fusion communities to explore leveraging existing knowledge technologies to benefit autonomous systems. The symposium was the first of its kind and was attended by 36 participants who represented a cross-section of the communities mentioned above.

Many researchers feel that an autonomous system must have an internal representation of entities, events, and situations that it perceived in the world in order for it to behave appropriately in uncertain environments. The term *autonomous systems* in this context refers to embodied intelligent systems that can operate for extended periods of time without human supervision. A major challenge for these systems is maintaining an accurate internal representation of pertinent information about the environment.

A large body of work exists in various knowledge representation, ontology, and data fusion areas, yet relatively little has been applied to real-time world modeling in autonomous systems. The primary goals of this symposium were threefold: first, to educate the autonomous systems community as to the strengths and weaknesses of various knowledge-representation approaches; second, to educate the knowledge-representation community as to the knowledge-related challenges being faced within the autonomous-systems arena; and third, to establish networks of teaming arrangements and possible collaborations to allow the communities to work closer together in the future. All three goals were accomplished to various levels of success, with the second and the third goals resulting in a bit more success than the first.

The symposium started with an autonomous systems keynote presentation by Ernst Dickmanns (University of the Federal Armed Forces of Germany, Munich) who described the state of the art in autonomous vehicle research and development, focusing on efforts over the past 25 years at his university. This was followed by six paper presentations that were grouped into the tracks "Knowledge-Representation Perspectives and Integration Issues" and "Knowledge Representation for Autonomous Mobility." Following the paper presentations was a poster session.

On the second day, Michael Genesereth (Stanford University) gave the knowledge-representation keynote presentation, titled "World Models for Autonomous Systems," in which he described his thoughts on what types of knowledge representations appeared to provide the most value to autonomous systems. This was followed by four paper presentations in the track "Applying Ontologies to Autonomous Systems."

Some questions seemed to be common following many of the presentations. They were "How does one know what knowledge should be embedded in an external knowledge base versus in the code itself?" "How does one know which representations are good for what types of requirements?" and "What is an ontology—how is it different from the knowledge representation techniques we have used in the past?" The answers to the first two questions varied from presenter to presenter, showing that there is no clear-cut answer and that more research needs to be performed. To address the third question, Michael Uschold from Boeing gave a brief, impromptu presentation at the start of the day describing a common view of what an ontology is and how it is intended to be used.

After the presentations, the audience was split into three predefined, cross-disciplinary breakout groups, each tasked with addressing a challenge problem. Their job was to determine a "knowledge architecture" for a group of five trash-removal robots that were responsible for cleaning an airport. The robots had to coordinate with each other, provide complete trash-removal coverage of the airport multiple times each day, monitor their health, travel within marked lanes whenever possible, recycle, identify suspicious packages, and stay a predefined distance from humans at all times. Within the "knowledge architecture," the groups had to define the types of knowledge necessary for the robots to perform their tasks, identify the types of representations that lent themselves best to representing that type of information, and develop the interfaces between the knowledge sources and the algorithms that were controlling the robots.

On the third day, the moderators of the breakout groups reported back on their groups' findings. Each group tackled a different aspect of the problem, often employing different approaches to do so, thus providing insight that there is no magic bullet in knowledge representations and that different techniques offer different advantages and disadvantages. Allowing the participants to get their hands dirty by addressing the challenge problem also confirmed the belief that knowledge representation for autonomous systems is a tough problem and should receive more attention from the community.

This was followed by a panel discussion made up of participants from industry and academia and represented all of the communities present at the symposium. The panel was tasked with highlighting the main issues and challenges that came out of the symposium, as well as determining the best way for these communities to work together in the future. Issues that arose from the panel included the need for an upper ontology, the challenge of integrating disparate terminology and semantics from different disciplines, and the need for a knowledge-representation formalism to capture the autonomous systems' competencies. There was also widespread agreement that the symposium was valuable and that similar ones should be held in the future.

The papers presented at this symposium are available as AAAI Technical Report SS-04-04.

– Craig Schlenoff
National Institute of Standards and Technolog

Language Learning: An Interdisciplinary Perspective

Language learning is a challenging problem for artificial intelligence. It encompasses concept development and perceptual development, social learning and imitation, as well as learning the lexicon, the grammar, and other aspects of language. It drives new technologies that apply widely to other kinds of sequential data. And because most of the world's knowledge is represented linguistically, machines are limited by their inability to understand language.

This symposium brought together representatives of several communities—the corpus-based and grounded language learning communities and the developmental psycholinguistics and language education communities—to assess progress in machine language learning and how what we know about human linguistic development might speed that progress. Linguists and psychologists were particularly well represented at this symposium.

Three kinds of interdisciplinary discussions took place. In grounded language learning, language describes a present scene and is often learned in a language game of some sort with a competent language user. Corpus-based approaches work with corpora of language dissociated from a present scene and not generated in a language game that includes the learner. Learning rates may be higher for grounded language learning; corpus-based approaches may learn a wider range of word classes, including words with abstract semantics that do not refer to a present scene.

Both approaches are inherently statistical, and it became clear that much can and should be shared between the practitioners of each. A second topic of discussion was the integration between lexical acquisition and grammatical inference. Knowing word meanings can help one acquire grammatical rules, and the assignment of words to grammatical categories should help acquire their meanings. A third topic of discussion occurred between language learning researchers and those who work on large, commonsense knowledge bases (such as Cyc). Language is layered on a conceptual system and depends on that system for its interpretation, and language conveys new concepts and distinctions. so language learning both depends on and extends commonsense knowledge.

The problem of learning syntax received a large amount of attention during the symposium. Presentations on this topic covered such diverse areas as inference of formal languages, corpus-based approaches to learning syntactic structure, unsupervised learning of context-sensitive languages, bootstrapping syntax from semantics, learning semantic parsers, and the role of errors in learning universal grammar. The discussions surrounding this topic benefited immensely from the interdisciplinary backgrounds of the attendees.

At the end of the symposium it was clear to all in attendance that much exciting work lies ahead and that the rate of progress can be greatly increased by intense interaction among the various relevant research communities.

The papers presented at this symposium are available as AAAI Technical Report SS-04-05.

– Tim Oates
University of Maryland

Semantic Web Services

While the semantic Web may be a vision of a new architecture for the World Wide Web characterized by the association of machine-accessible formal semantics, the development of representational issues and logical frameworks (such as the Web ontology language, OWL) will take us only so far. To fully realize this vision, behavioral issues (such as interactions between semantic Web services) will have to be tackled. Serendipitous interoperability, that is, the unarchitected, unanticipated encounters of agents on the Web is an important component in this realization.

Nearly 50 researchers gathered at the Semantic Web Services Symposium to discuss the challenges of describing, composing, invoking, and utilizing semantic Web services, including the pragmatic issues of mapping between different ontological representations, managing security, and defining and utilizing access policies. From a total of 51 paper submissions, 20 were selected and presented at the conference, with two invited talks, a panel session, and a breakout session.

Service description and discovery support the provision of dynamic service deployment within an open environment. With no a priori knowledge of the diversity of services available, a need to express requirements that can be then matched against advertised service descriptions was identified. A variety of solutions were proposed that included negotiation, brokerage, and structured queries. The service-oriented approach facilitates the construction (or composition) of new services based on existing (and consequently discovered) services. Semantic reasoning can be used to resolve many of the mismatches between heterogeneous services, and consequently, the majority of papers focused on this area, from the perspective of planning, workflow construction, model checking, concept-based constraint satisfaction, and interactive service composition. Security is a major concern, however. How can one trust that a service will actually adhere to its advertised capability and not act inappropriately (such as distribute personal information)? Likewise, services may be distributed across many different servers within different institutions, and hence could be liable to malicious attacks (such as denial of service attacks). One session focused on this area by examining how security measures could be encoded within service descriptions, how policy languages could be used to determine appropriate behavior, and how a semantic fire-wall architecture utilizes these to protect an institution's services.

Not the entire symposium concentrated on the mechanics of semantic reasoning for services. Several ontologies that support service provision were presented, from defining time and duration to upper-level ontologies that augment the expressivity of existing frameworks, such as OWL-S and UPML, as well as approaches to align different representations or even learn new representations based on existing repositories.

The invited talks illustrated both the need and utility of semantic Web services within knowledge-rich domains (AKT) and to support large-scale Grid computing within emerging eScience. Emerging standards work was presented by representatives from the Semantic Web Services Initiative, and a breakout session allowed participants to discuss pressing issues regarding OWL-S, the use of semantics with the Grid, challenges in both composition and invocation of services, and augmenting service descriptions with additional metadata that support and enhance the capability description, thus improving discovery.

By the end of the symposium, it was clear that semantic Web services represented a fusion of diverse AI research areas such as agent-based systems, description logics, computational grid, and Web services. Though many challenges still remain, this emerging field will benefit from the synergies emerging in this area.

The papers presented at this symposium are available as AAAI Technical Report SS-04-06.

– Terry Payne
University of Southampton