

Reports on the 2004 AAAI Fall Symposia

Nick Cassimatis, Sean Luke, Simon D. Levy, Ross Gayler, Pentti Kanerva, Chris Eliasmith, Timothy Bickmore, Alan C. Schultz, Randall Davis, James Landay, Rob Miller, Eric Saund, Tom Stahovich, Michael Littman, Satinder Singh, Shlomo Argamon, and Shlomo Dubnov

- The American Association for Artificial Intelligence presented its 2004 Fall Symposium Series Friday through Sunday, October 22–24 at the Hyatt Regency Crystal City in Arlington, Virginia, adjacent to Washington, DC. The symposium series was preceded by a one-day AI funding seminar. The topics of the eight symposia in the 2004 Fall Symposia Series were: (1) Achieving Human-Level Intelligence through Integrated Systems and Research; (2) Artificial Multiagent Learning; (3) Compositional Connectionism in Cognitive Science; (4) Dialogue Systems for Health Communications; (5) The Intersection of Cognitive Science and Robotics: From Interfaces to Intelligence; (6) Making Pen-Based Interaction Intelligent and Natural; (7) Real-Life Reinforcement Learning; and (8) Style and Meaning in Language, Art, Music, and Design.

The American Association for Artificial Intelligence presented its 2004 Fall Symposium Series Friday through Sunday, October 22–24 at the Hyatt Regency Crystal City in Arlington, Virginia, adjacent to Washington, DC. The symposium series was preceded on Thursday, October 21 by a one-day AI funding seminar, which was open to all registered attendees. The titles of the eight symposia in the 2004 AAAI Fall Symposia Series were: (1) Achieving Human-Level Intelligence through Integrated Systems and

Research; (2) Artificial Multiagent Learning; (3) Compositional Connectionism in Cognitive Science; (4) Dialogue Systems for Health Communications; (5) The Intersection of Cognitive Science and Robotics: From Interfaces to Intelligence Making; (6) Pen-Based Interaction Intelligent and Natural; (7) Real-Life Reinforcement Learning; and (8) Style and Meaning in Language, Art, Music, and Design. Reports on each of these symposia are included in this report.

The highlights of each symposium were presented at a special plenary session. Notes were prepared and distributed to participants in each symposium, and papers from all but one of the symposia (Real-Life Reinforcement Learning) are also available as AAAI Technical Reports.

The AI Funding Workshop

The AI Funding Workshop, held on Thursday, October 21, 2004, provided an opportunity for new and junior researchers—as well as students and postdoctoral fellows—to get an inside look at what funding agencies expect in proposals from prospective grantees. Representatives and program managers from various funding agencies such as the Defense Advanced Research Projects Agency (DARPA), the National Science Foundation (NSF), the Office of Naval Research (ONR),

and the Department of Homeland Security gave presentations. Several successful researchers discussed what they believed made them successful, and provided advice on how to play the funding game.

Achieving Human-Level Intelligence through Integrated Systems and Research

The AAAI Fall Symposium on Achieving Human-Level Intelligence through Integrated Systems and Research was motivated by the belief that increasing subfield specialization within the AI community was causing the field to lose sight of one of its original goals—achieving human-level artificial intelligence—and that combining the insights of subfields into integrated systems and research programs could renew progress towards this important goal. Although these themes date back at least to the work of Allan Newell, recent progress and promising new approaches merited a symposium on the topic.

Many participants presented work that either developed or was based on architectures whose purpose was in part to enable the design of integrated systems. Two insights that came from this part of the symposium were that (1) many underlying similarities and points of contact exist between different architectures that could potentially enable work based on each one to be integrated and (2) architectures needed to incorporate multiple representational formalisms and inference methods to achieve human-level AI.

Symposium attendees repeatedly discussed the challenge of evaluating integrated intelligent systems. There was consensus among participants that metrics in machine learning, planning, and natural language processing have driven advances in those subfields, but that those metrics have also distracted attention from how to combine their methods to create systems with human-level intelligence. Several ideas for metrics to evaluate and motivate human-level AI work—including tests of mental ability (IQ

tests, tests of creativity, tests of mechanical ability, and so on)—were discussed. Domains for motivating, testing, and funding this research were proposed (some during our joint session with the Human-Robot and Computer Interaction symposium), including autonomous vehicle navigation, synthetic characters for training, and companion cognitive systems.

The participants proposed several steps for advancing towards human-level intelligence. These steps included creating better methods of evaluating progress; developing methods for bootstrapping from existing intelligent systems to greater intelligence; compiling lists of qualitative descriptions of problems that people solve; identifying and implementing common reasoning, planning, and learning abilities required by many domains; identifying the neural circuitry that underlies human intelligence; and tackling problems involving a social dimension that forces systems to behave in a plausibly human manner.

Nick Cassimatis, *Rensselaer Polytechnic Institute*

Artificial Multiagent Learning

Multiagent learning as a community is spread over fields as diverse as robotics, evolutionary computation, and networking. The goal of this symposium was to bring together multiagent learning researchers from these and other fields to discuss common issues and interests. Multiagent systems are the part of distributed AI that emphasizes the joint behaviors of at least partly autonomous agents, in environments where constraints limit the degree to which such agents may know about the entire world state (including one another). We strove to make the topic of multiagent learning as inclusive as possible: the application of machine learning to multiagent systems (whether with a single learner or with many).

Interest in multiagent learning has grown rapidly in the last few years as networking, simulation, software agent, and robotic technologies have become more easily accessible. Still, multiagent learning faces daunting

game-theoretic and scalability challenges stemming from increasingly large numbers of agents, more complex agent behaviors, partially observable environments, and mutual adaptation.

Seventeen papers were presented on topics including reinforcement learning, evolutionary computation, multiagent simulation, game theory, distributed constraint satisfaction, network analysis, and artificial life. In addition to papers, the symposium hosted discussion of problem domains and literature surveys, plus demonstration of multiagent simulation systems and development environments. The symposium cosponsored a joint session and panel discussion with the Real-Life Reinforcement Learning Symposium.

The lion's share of the symposium papers focused on four major topics. Foremost was analysis of stochastic and repeated games with multiple learners. This is a common topic in multiagent learning literature, and it is fraught with game-theoretic challenges due to the presence of multiple learners "moving the goalposts" on one another. Much of this topic deals with guarantees of convergence to Nash equilibria, and this singular focus bore harsh criticism from one particular paper by Yoav Shoham, Rob Powers, and Trond Grenager (Stanford University). Another major topic was the evolutionary discovery of efficient cooperators through the use of tags. Here the agents may choose whom to cooperate with, and evolve into cliques that work effectively with one another.

A number of papers dealt with evolutionary and coevolutionary optimization of teams of agents whose behaviors (such as robotic predator functions) cannot easily be represented as strategies for games. There are strong theoretical relationships between this area and the repeated game literature. The paper by Elena Popovici and Kenneth De Jong (George Mason University) attempted to explain visually the convergence properties—or lack thereof—of coevolutionary optimization.

Finally, several papers presented work in adaptive agents in a dynamic network of communication. This area brings together social network analy-

sis, constraint satisfaction, and graph theory. The topic is of increasing interest with the advent of peer-to-peer network services and with ad-hoc wireless networks among mobile agents.

This symposium continues a tradition of multiagent learning symposia that began with the 2002 AAAI Spring Symposium on Collaborative Learning Agents. A future opportunity to continue the discussion will come at the AAAI-05 workshop on multiagent learning this coming July, 2005.

Sean Luke, *George Mason University*

Compositional Connectionism in Cognitive Science

Compositionality (the ability to combine constituents recursively) is generally taken to be essential to the open-ended productivity of perception, cognition, language, and other human capabilities aspired to by AI. Ultimately, the neural networks of the brain implement these capabilities, yet connectionist models have had difficulties with compositionality. This symposium brought together connectionist and nonconnectionist researchers to discuss and debate compositionality and connectionism.

The aim of the symposium was to expose connectionist researchers to the broadest possible range of conceptions of composition—including those conceptions that pose the greatest challenge for connectionism—while simultaneously alerting other AI and cognitive science researchers to the range of possibilities for connectionist implementation of composition. We therefore welcomed and encouraged submissions from both proponents and critics of connectionist representations, as long as the work described focuses on compositionality in the context of AI or cognitive science.

The symposium was characterized by a wide variety of approaches to the compositionality issues—including holographic reduced representations, Hebbian learning, fractals, and linear logic—and by vigorous debate on both the fundamental questions and the merits of these approaches. A point of agreement that emerged from these

discussions was the significant difference between (1) research concerned with the compositional nature of cognitive representations useful to AI and other practical domains, and (2) research whose primary focus is modeling specific empirically observed patterns in human language. The latter approach was championed by Jeff Elman (University of San Diego), our first plenary speaker, whose talk described the surprising level of grammatical richness that can be learned by his simple recurrent network model. The two approaches were unified in the second plenary talk by Mark Steedman (University of Edinburgh), who described how such linguistic patterns might have emerged from the cognitive constraints imposed by planning tasks in primates.

Other highlights of the meeting included a panel discussion on the relative merits of vector-based methods against temporal-synchrony methods in solving the notorious variable-binding problem, led by Tony Plate and John Hummel (University of California, Los Angeles), and a related presentation of recent work on that topic by Ross Gayler (La Trobe University).

Simon D. Levy, *Washington and Lee University*; Ross Gayler, *La Trobe University*; Pentti Kanerva, *Redwood Neuroscience Institute*; Chris Eliasmith, *University of Waterloo*

Dialogue Systems for Health Communication

There is a growing body of research over the last two decades in both artificial intelligence (particularly in automated dialogue systems) and healthcare (including health communication, health behavior change, and medical informatics) focusing on the development of automated systems that interact directly with patients to achieve health education and behavior change. This symposium was a first attempt to bring these two research communities together. This area of research will have a potentially large impact on society, given that it can address health-behavior problems such as smoking, poor diet, lack of exercise, and medication nonadherence—problems that have been blamed for up to

60 percent of the \$1.2 trillion per year spent on healthcare in the United States.

Twenty-six participants attended the symposium. They came from industry, government, and academe. The eleven paper presentations ranged from descriptions of fielded systems using a variety of interactive media (PDAs, telephones, wearable devices, desktop computers, speech, embodied conversational agents, and interactive computer games) to general methods for dialogue planning, representation and text generation of uncertain information, and system evaluation. Applications spanned health behavior change, chronic disease management, social support for caregivers, genetics counseling, administration of informed consent, emergency room triage, medication adherence, and real-time speech-based machine translation for physicians. Papers were also presented on methods for diagnosing medical and mental health conditions based on speech and written language samples.

Invited talks were given by Robert Friedman, M.D., a professor of medicine at the Boston University School of Medicine, and Geoffrey Clapp, CTO of HealthHero Network. Both speakers have experience fielding health dialogue systems and conducting clinical trials involving thousands of patients. The results of these trials have shown that the systems are efficacious, and they have been used for a wide range of applications in health behavior change and chronic disease management.

One panel discussion covered the topic of whether health dialogue systems constitute a unique field of inquiry or just an application area. Several unique aspects of these systems were identified, including the special needs for privacy and security; reliability; long-term use (up to a lifetime for chronic disease management); continuity and persistence of dialogue state between uses; negotiation of treatment regimens; and establishing and maintaining trust and social bonds between automated healthcare providers and patients. A second panel discussed tools, scripts and corpora that could be shared by researchers working in this

area. A brainstorming session was also spent on the challenge problem of how health dialogue systems can motivate users to stay engaged with them over months or years of repeated use. Many ideas were generated, ranging from the humorous (involving addictive substances) to the theoretical.

Because of the strong attendance at the symposium and the high interest by participants there are plans to continue activities in this area at future AAAI fall symposia and other venues. Information on the symposium and related information can be found at the symposium Web site.¹

Timothy Bickmore, *Boston University School of Medicine*

The Intersection of Cognitive Science and Robotics: From Interfaces to Intelligence

Principles and methodologies from cognitive science are beginning to be applied to autonomous robots. The use of cognitive science in robotics takes varied forms, from using computational cognitive models as reasoning mechanisms for robots, to the design and control of human-robot interaction (HRI). This interdisciplinary symposium brought together researchers in robotics, cognitive science, and human-machine interfaces, as well as philosophy, mathematics, and ethics, to examine this emerging area.

Because this was an interdisciplinary symposium, we set definitions ahead of time. We focused on cognitive science work that had some cognitive plausibility (that is, it could arguably be claimed that the representation, strategies, and/or actions had some basis in human cognition) or “person in the-loop” issues. We also focused on robotics work that emphasized embodied systems, such as mobile robots and autonomous vehicles, not just software agents.

The symposium consisted of short and long presentations, extended discussion periods, and joint sessions with other symposia. Presented work covered a wide range of topics, includ-

ing (1) the use of computational cognitive models as reasoning mechanisms in robots and as a model of humans or other agents with whom the robot must interact; (2) models of emotion; (3) human subject studies to determine effective interaction models for a given role and task domain; (4) development of cognitively plausible memory models and their affect in several applications; (5) human-robot interaction studies; and (6) using cognitive methods for visual processing and understanding.

Several themes emerged in many of the talks. Higher-level cognitive compatibility is important for HRI because it results in reduced cognitive load, thus making the robot actions and intentions easier to understand (and therefore making it easier to recognize and fix misunderstandings), and also making the robot more predictable (and therefore helping to engender trust, among other benefits). Inferring intents, goals, and motivations, taking visual perspective, imitation learning, and spatial reasoning were all identified as important for human-robot collaboration, communication, and learning.

Two of our sessions were held with other symposia. In the early afternoon on the second day, we joined the Real Life Reinforcement Learning symposium. In the late afternoon, participants of the Achieving Human-Level Intelligence through Integrated Systems and Research symposium joined us. Each symposium supplied a speaker and allowed time for discussion.

On the final day we held a session where we thought about where cognition in robots might lead in the far future, and also had a spirited philosophical discussion of safety and ethics.

Alan C. Schultz, *Naval Research Laboratory*

Making Pen-Based Interaction Intelligent and Natural

As personal digital assistants (PDAs) and tablet computers grow in commercial popularity, researchers in pen-based computing are faced with the challenge of moving these interfaces

forward. We need to look beyond handwriting recognition, and beyond simple ink note-taking, to learn what it takes to create systems whose interfaces are as fluid and natural as drawing on paper, and yet invested with the power of symbolic computation. How can freehand sketches and diagrams be interpreted? How can domain knowledge be represented and used in sketch interpretation? How can we incorporate diverse streams of information about sketches—spatial, temporal, video, audio—to resolve ambiguity and achieve greater accuracy? What novel user interfaces can be driven by sketching?

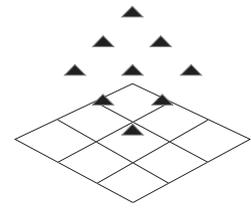
These questions and others drew thirty-two researchers, a cross-section of the artificial intelligence and human-computer interaction (HCI) communities from both academe and industry, to meet at this symposium.

Invited speaker Gordon Kurtenbach (Alias Systems) delivered an entertaining and enlightening talk that challenged the symposium to reconsider what we mean by “natural.” Using examples from a decade of research on user interfaces, both pen-based and otherwise, Kurtenbach argued that naturalness depends strongly on context, culture, and expertise.

The twenty-five papers presented at the symposium ran the gamut from low-level recognition to high-level applications. Issues discussed fell roughly into five areas: (1) early stages of recognition, in which raw ink strokes are segmented and interpreted as line segments, arcs, and other simple geometric shapes; (2) late-stage recognition, in which shapes are combined into symbols and relationships are discovered using domain knowledge; (3) sketches of three-dimensional objects or inside three-dimensional environments; (4) applications, including animation, lecture annotation, and sketch beautification; and (5) user interface issues, including design guidelines and techniques for reducing mode errors in sketching interfaces.

Demonstration sessions, held just before dinner each day, were the highlights of the symposium, offering participants a chance to see new systems in action.

A working lunch session on the last



*Please Join Us
for the 2005
Fall
Symposium
Series*

November 3–6, 2005

Hyatt Regency Crystal City

Arlington, Virginia

www.aaai.org/Symposia/Fall/2005/

day focused on a discussion of research issues and possible “killer” applications for pen-based computing. Among the issues identified were how to manage the tradeoffs of domain dependence; how to design sketching interfaces that are easy for novice users, but transition smoothly to efficient and powerful use by experts; and how to develop more infrastructure for the pen-based research community, including not only corpora but also standard interfaces, prototyping tools, and evaluation methods.

Randall Davis, *Massachusetts Institute of Technology*; James Landay, *University of Washington*; Rob Miller, *Massachusetts Institute of Technology*; Eric Saund, *Palo Alto Research Center*; Tom Stahovich, *University of California, Riverside*

Real-Life Reinforcement Learning

The Real-Life Reinforcement Learning (RLRL) symposium brought together researchers interested in developing intelligent systems that make decisions by attempting to maximize some measurable performance objective. The emphasis of the symposium was on discussing algorithms that are ef-

fective when learning from “real-life” (measured) data. The goal was to build a community by highlighting existing RLRL efforts and sharing experiences and approaches that have been demonstrated to be successful.

The meeting attracted approximately fifty participants from five countries. We were gratified to hear from an insightful and enthusiastic collection of researchers including a few reinforcement-learning pioneers and even several reluctant “converts” who discovered the utility of reinforcement-learning approaches when other techniques had failed them.

As the focus of the meeting was on learning systems that interact with real life environments, many of our presentations concerned robotic applications. Learning in a robotics setting is especially challenging due to expensive exploration and noisy sensors. A wide variety of robotic subproblems were described including motion planning, navigation, grasping and manipulation, and active-sensing strategies. The robotic platforms discussed included Lego Mindstorms, Sony Aibo legged robots, an air-hockey-playing torso, and a baseball bat swinging arm. The presenters advocated robotic examples as among the most exciting domains for studying RLRL.

Other interesting problem domains that were also described included control in detailed simulations (automotive, rocketry, and computer games), computer-system resource allocation, financial trading, and process control. While the expense of exploration in these domains is often less than in robotic domains, the sheer combinatorics of some of these problems made them just as daunting.

In addition to focused presentations, we undertook a series of more open-ended discussions on topics ranging from the AI problem, multiagent learning, hierarchical representations, to cognitive robotics. We also participated in two shared sessions with the Artificial Multi-Agent Learning and The Intersection of Cognitive Science and Robotics symposia, and benefited from the unique perspectives these communities provided.

As the meeting wrapped up, the participants expressed a willingness to

continue our community-building efforts. Ongoing projects include creating a repository of real-life RL examples, assembling a web page debunking harmful myths about reinforcement learning, and planning a journal special issue on the topic. The papers from this symposium were not published.

Michael Littman, *Rutgers University*;
Satinder Singh, *University of Michigan*

Style and Meaning in Language, Art, Music, and Design

This highly multidisciplinary symposium brought together researchers working with computational models for style in a great variety of domains, including the visual arts, music, natural language, theater and cinema, game playing, and architecture. During very stimulating meetings we explored many different perspectives on style and its role in various facets of human behavior, attempting to create an integrated research community from our heretofore disparate research enterprises.

There was broad agreement that style is generally expressed by a confluence of a great many small surface details in any particular work (painting, concerto, essay), which collectively point towards a particular coherent interpretation of a feeling or identity. This distributed character of stylistic expression was commented upon by several speakers, and was contrasted with the compositionality of classical denotational accounts of meaning. But style is far more than just a collection of features, and a deeper understanding must also take into account the history and social context of a work. That is, production constraints, whether externally or internally imposed, are key to understanding how features may cohere into a recognizable style. Also relevant was the observation of how stylistic patterns may be fundamentally altered by perceivers changing point of view, underscoring the essentially context-sensitive nature of stylistic apperception.

At a high level, three main perspectives as to style’s place in cognitive

processing were discussed. One view of style had it emerging from an interaction between production constraints and a coherent (though possibly initially vague) artistic vision. An opposed view of style saw it as a realization of explicit affective and interpersonal goals, via many small features distributed in a work. A third notion of style was a consistent strategic stance enabling decision-making when various alternatives seem equally good from a utility-theoretic perspective. More specific approaches to analyzing stylistic features of works in various media were discussed, including the role of metaphor in stylistic expression, anticipation and familiarity in emotional response, rules of style in construction of a stylistic ideal, and functional and rhetorical organization in constructing a style for a community of discourse.

It was clear that whatever it really is, style is an essential part of human intelligence, and that a computational understanding of style becomes more important as our lives become intertwined with a ubiquitous information load and ever-more capable information technologies. In terms of applications, we all agreed that style research will be increasingly central in building bridges of understanding between people and intelligent systems in areas such as machine translation, user-adaptive interfaces, or intelligent agents. Beyond applications, though, an understanding of style, from the mundane (peoples’ walking styles, for example) to the sublime (Monet’s lilies, or Beethoven’s symphonies), is fundamental to understanding what it means to be human. A new community is coalescing to study this age-old problem, informed by the diverse traditions of artificial intelligence, cognitive science, and psychology, together with humanities scholarship, artistic expression, and philosophical inquiry.

Shlomo Argamon, *Illinois Institute of Technology*; Shlomo Dubnov, *University of California at San Diego*

Note

1. www.misu.bmc.org/~bickmore/dshc/