The 1995 Fall Symposia Series

Sponsored by the American Association for Artificial Intelligence

■ The American Association for Artificial Intelligence (AAAI) held its 1995 Fall Symposia Series on 10 to 12 November in Cambridge, Massachusetts. This article contains summaries of the eight symposia that were conducted: (1) Active Learning; (2) Adaptation of Knowledge for Reuse; (3) AI Applications in Knowledge Navigation and Retrieval; (4) Computational Models for Integrating Language and Vision; (5) Embodied Language and Action Symposium; (6) Formalizing Context; (7) Genetic Programming; and (8) Rational Agency: Concepts, Theories, Models, and Applications.

The American Association for Artificial Intelligence (AAAI) held its 1995 Fall Symposia Series on 10 to 12 November in Cambridge, Massachusetts. This article contains summaries of the eight symposia that were conducted: (1) Active Learning; (2) Adaptation of Knowledge for Reuse; (3) AI Applications in Knowledge Navigation and Retrieval; (4) Computational Models for Integrating Language and Vision; (5) Embodied Language and Action Symposium; (6) Formalizing Context; (7) Genetic Programming; and (8) Rational Agency: Concepts, Theories, Models, and Applications.

Active Learning Symposium

Many learning algorithms are passive in that the learner simply observes its environment (for example, by accepting random training samples). In contrast, active learners have the ability to influence their environment and, thus, the training data they receive. The Symposium on Active Learning examined these active-learning approaches, bringing together researchers from such areas as machine learning, statistics, robotics, control, game theory, and machine reasoning.

The actions available to a learner include choosing what data to remember (and forget), asking questions of a teacher, roaming around an environment, and planning and executing complex experiments. Not surprisingly, no single algorithm or theory covers all these situations, but common themes emerge nonetheless. Most approaches assume that learning will be sequential, so that the results of earlier actions guide the selection of later ones. Many approaches model a learner's uncertainty in various parts of a domain and use this information to guide learning. Similarly, many approaches model what the learner needs to know, whether this modeling is explicit (for example, goals) or implicit (for example, a distribution of examples that the system is likely to be tested on).

The American Association for Artificial Intelligence (AAAI) held its 1995 Fall Symposia Series on 10 to 12 November in Cambridge, Massachusetts. This article contains summaries of the eight symposia that were conducted.

Of course, active learning shares many concerns with other learning research, including the role of prior knowledge and bias, the tension between theoretical neatness and computational tractability, and the centrality of representation. Another commonality is the wide applicability of the research. Participants in the symposium found active learning applicable to such diverse areas as information retrieval, robotic control, computer security, oncology, and weather modeling.

David Cohn
Massachusetts Institute of Technology

David Lewis
AT&T

Adaptation of Knowledge for Reuse

The Symposium on Adaptation of Knowledge for Reuse, inspired by the need for adaptation theories with practical implications for case-based reasoning (CBR) and other AI tools, provided a forum where attendees with a wide range of interests shared ideas on adaptation and reuse.

Invited speaker David Leake outlined the CBR perspective, and speakers Mark Keane (analogy), Maja Mataric (multiagent systems), and Larry Latour (software reuse) presented complementary perspectives. Ralph Barletta explained why automated adaptation strategies do not exist in current CBR applications, and Kevin Ashley presented reasons for more optimism on their future use.

Paper presentations targeted these and several other areas (for example, knowledge acquisition, plan reuse, theorem proving) on topics including reuse architectures, hierarchical CBR, case-based adaptation, higher-order extensions to explanation-

Discussions were spirited and entertaining. It is heartening to note that although the roboticists and computational linguists often disagreed, there was little or no tendency for one community to marginate the other. Each community readily accepted the other's issues as hard but important problems, and each community was generally interested in learning about the other field.

based generalization, target reformulation, and the constraint of adaptation using causal models. A poster session, panels, breakout sessions, and postpresentation discussions led by senior attendees allowed participants to communicate and refine their messages.

Discussion topics included benefits from integrating adaptation with other problem-solving phases; knowledge sources for adaptation; interactive adaptation in human-machine combinations; usefulness of reuse versus adaptation; diversity of existing adaptation strategies; adaptation of individual cases or solutions versus synthesis of multiple ones; the need for domain-dependent adaptation knowledge and techniques in many applications; the relationships between reformulation, adaptation, analogical mapping, and learning; and the exploitation of implicit knowledge to improve efficiency, problem-solving range, and solution quality.

A paper and an edited volume based on this symposium are being planned. Details can be found at http://www.aic.nrl.navy.mil/~aha/aaai95-fss.

David W. Aha
Naval Research Laboratory

AI Applications in Knowledge Navigation and Retrieval

There is an emerging recognition that the growth of the World Wide Web and other large-scale information systems is rendering existing information tools inadequate. This area has become fertile ground for AI research and the subject of several recent and future American Association for Artificial Intelligence symposia and workshops. The AI Applications in Knowledge Navigation and Retrieval Symposium provided a snapshot of current AI approaches to information systems, concentrating particularly on problems associated with navigating large, dynamic information sources.

In emphasizing navigation, the symposium looked particularly at user interface issues in complex information spaces. One of the functions of browsing is to give the user a feel for what is contained in the information space. Much of the research reported at the symposium looked at the problem of providing intelligent assistance to help users find what they want and still offering a sense of "what's out there."

The symposium emphasized implemented applications. Eighteen demonstrations were given at three hands-on demonstration sessions presented at the Massachusetts Institute of Technology AI Lab. These demonstrations and other presentations of existing systems helped stimulate concrete discussion, throughout the symposium, of problems and potential solutions. Problems that were addressed ranged from guided web browsing to the personalization of dedicated hypermedia document bases. Although the participants came from a wide range of backgrounds and institutions, there was a great deal of common ground among approaches. Many of the systems described at the symposium combined standard information-retrieval methods with AI techniques, including machine learning, natural language processing, knowledge representation, and case-based reasoning.

Robin Burke
University of Chicago

Computational Models for Integrating Language and Vision

Integrating language and vision has been a long-standing goal of AI researchers because much of human communication centers on events and activities in the visual world. The Symposium on Computational Models for Integrating Language and Vision began with an overview of the field, summarizing progress since Terry Winograd's SHRDLU in 1970. The observation was made that although true integration at the conceptual level might still be an elusive goal, much progress has been made in establishing useful correspondence between the two modalities to enable several exciting applications. The question of whether this topic constitutes a new field in AI was posed; because both the problems and the solutions are novel, the consensus was yes.

There were several paper sessions covering such topics as event perception, gesture understanding, knowledge representation, spatial reasoning, diagram-image understanding, and image-video indexing. Joe Mundy of GE gave an invited talk

entitled "Image Observation Events: A Link to Linguistic Context." In it, he summarized the current trend in computer vision toward three-dimensional model-supported image understanding, multiple sensors, and so on. He concluded by outlining the key role that language could play in such an environment.

The panel and discussion sessions addressed issues such as the need to (1) tackle easier, lucrative applications (low-hanging fruit); (2) develop integrated language-vision ontologies; (3) focus on both virtual environments as a test bed in tackling deep issues in correspondence as well as immediate applications where shallow correspondence would suffice; and (4) define the scope of the field (that is, exclude pure vision or natural language understanding tasks). It was decided not to try and impose artificial boundaries around a newly emerging field. Finally, there was a general desire to hold a follow-up conference in two years.

Rohini K. Srihari
State University of New York at Buffalo

Embodied Language and Action

The Symposium on Embodied Language and Action focused on agents that use language and gesture to facilitate extended interactions in a shared physical or simulated world. The attendees consisted mostly of computational linguists and roboticists. The goal was to discuss how embodiment in a shared world stimulates communication and provides a resource for understanding it. One challenge was a lack of consensus on what the term *embodied language* actually meant. Interpretations ranged from treating language use as situated activity to grounding of lexical semantics in sensory-motor systems.

The symposium began with presentations of work by the participants, ranging from an implemented vision system that interprets pointing gestures to a robot cat with limited speech recognition to the collection of empirical data on human judgments of noun-phrase referents in the context of a specific task (*collaborative navigation*). Later, the attendees divided into groups to discuss specific task scenarios and how they might be solved by language-using agents. Although there was no way to develop true solutions, discussion led to the examination of important issues such as what subproblems were hard or easy or how the context of the task could be used to simplify these subproblems.

Discussions were spirited and entertaining. It is heartening to note that although the roboticists and computational linguists often disagreed, there was little or no tendency for one community to marginate the other. Each community readily accepted the other's issues as hard but important problems, and each community was generally interested in learning about the other field.

Ian Horswill
Massachusetts Institute of Technology

Formalizing Context

The intuitive notion of context is pervasive in our cognitive world. Thus, systems that provide knowledge-level support need the ability to reason about contexts. The Symposium on Formalizing Context brought together researchers interested in formalizing this intuitive notion of context and providing a basis for developing such systems.

The first step in this direction is to provide a formal language that enables us to explicitly talk about contexts. Over the past few years, a number of logics have been developed with this goal in mind. They treat contexts as formal objects, thereby enabling them to represent properties of, and arbitrary relations between, contexts.

To further develop the formalization, we need to identify salient features of contexts. This task is especially difficult because of the various roles that context can play, depending on the application in which it is used. For example, the role context will play in integrating databases will be different from its role in natural language-discourse processing. Furthermore, its role in this processing will range anywhere from resolving references of pronouns to capturing the intentions and beliefs of the discourse participants. An important task before us then is to compare and contrast the different notions of context to identify what concepts they share and the ways in which they differ. We have already identified and formalized some basic concepts that seem to be relevant to most notions of context, including (1) a proposition holding or being true in a context, (2) a term having a value specific to a context, and (3) one context being just like another except that it makes some additional assumption. The symposium helped shed light on these issues.

Sasa Buvac
Stanford University

Genetic Programming

Genetic programming extends the genetic algorithm to the domain of computer programs. In genetic programming, populations of programs are selectively bred to solve problems. Starting with a primordial soup of hundreds or thousands of randomly created programs composed of functions and terminals appropriate to the problem, the population is progressively evolved over a series of generations by applying the operations of Darwinian selection and crossover (sexual recombination).

The Symposium on Genetic Programming was the first symposium devoted to this rapidly growing field. The application areas of genetic programming discussed at the symposium included computer system intrusion detection, the control of a real robot, signal classification and understanding, function modeling, and game theory. Genetic programming has begun to improve on human performance in real-world domains (for example, protein classification) and continues to contribute to machine learning in benchmark domains, for example,

the discovery that the 2-pole balancing problem is linearly separable and the creation of a majority rule for 149-element cellular automata, which improves on the three best handwritten rules (and all learned rules).

Genetic programming contributes to the field of machine learning by applying the crossover operation (common across all types of genetic algorithm) to dynamic representations (computer programs) that have position-independent constituents (computer operations). Under these circumstances, computational mechanisms specialized for a problem domain can emerge. This area of research is relatively new, and most of the symposium was spent pinpointing and fleshing out the salient issues that arise when investigating this complex system. Our discussion topics included the automatic hierarchical breakdown of a problem by way of evolved modular programs (for example, automatically defined functions); representations for evolved programs (for example, the ontogeny that results from implementing a growth phase before evaluating fitness); the interactive coevolution of subpopulations; and evolving programs that read from, and write to, memory. David Cohn, a guest speaker from the Active Learning Symposium, led a discussion on integrating active-learning techniques in genetic programming.

Short presentations paralleled the 19 working note papers, and the open discussion was energized and enlightening. It has been clear since the Fifth International Conference on Genetic Algorithms in 1993, where "the most enthusiastic subgroup at the conference was composed of those interested in genetic programming" (M. de la Maza, *AI Magazine* 15[2]: 83–85), that the open-discussion format of AAAI symposia is perfectly conducive for a meeting of genetic programming researchers. In addition to the working notes (available through AAAI), we have compiled an archive of our collective brainstorming that can be found at http://www.cs.columbia.edu/~evs/gpsym95.html.

The first genetic programming conference will be on 28 to 31 July in Stanford, California. For information, see http://www.cs.brandeis.edu/~zippy/gp-96.html. To join our online discussion of genetic programming, send mail to genetic-programming-REQUEST@cs.stanford.edu.

Eric V. Siegel
Columbia University

Rational Agency: Concepts, Theories, Models, and Applications

The Symposium on Rational Agency explored rational agency concepts and their implications for theory, research, and practice. The view that intelligent systems are, or should be, rational agents underlies much of the work in AI and cognitive science. However, no consensus exists on rationality standards for agents or even on the concept of agency itself.

Although our symposium reached no consensus, perhaps its chief contribution was to facilitate a constructive exchange of views among an international group of scholars drawn from an exceptionally diverse range of disciplines. Our participants represented not only AI, distributed AI, philosophy, and linguistics but also psychology, economics, and even organization science. By devoting significant symposium time to discussion, participants were able to explore some challenging and informative notions.

The discussion of agency mirrored AI's ongoing debate over agents as purposive and deliberative versus agents as fundamentally reactive. The contrast between agent models based on (defeasible) logical deduction and those founded on decision-theoretic choice processes was also well represented. However, another, intriguing issue consumed a good deal of our attention—the notion that adaptivity is fundamental to agency.

These perspectives on agency arose mainly within discussions of rationality principles that can be framed in terms of the following contrasts:

The discussion of agency mirrored AI's ongoing debate over agents as purposive and deliberative versus agents as fundamentally reactive.

(1) *perfect* (unbounded) *rationality* versus resource- and capacity-limited rationality as a suitable standard; (2) *epistemic rationality*, rationality principles for belief management, in constrast to *strategic rationality,* principles for rational (choice of) action; (3) the *prospective rationality* of deliberation versus some standard for the *retrospective rationality* of adaptation or learning; and (4) *social rationality* standards in contrast, or in addition, to individual rationality criteria.

Exploring social and individual rationality revealed two further contrasts: (1) rationality as *competence,* some criterion for the internal structure or function of the agent or social system, versus rationality as *coherence*, some external standard of performance or structure, and (2) social rationality as an a priori standard versus social rationality as emergent from the interactions among certain types of agent systems. These contrasts helped frame and facilitate our symposium's exploration of rational agency. They should be similarly useful to discussion within the intellectual community at large.

Michael Fehling
Stanford University