The 1996 Fall Symposium Series

Sponsored by the American Association for Artificial Intelligence

■ The American Association for Artificial Intelligence (AAAI) held its 1996 Fall Symposia Series on 9 to 11 November in Cambridge, Massachusetts. This article contains summaries of the seven symposia that were conducted: (1) Configuration; (2) Developing Assistive Technology for People with Disabilities; (3) Embodied Cognition and Action; (4) Flexible Computation: Results, Issues, and Opportunities; (5) Knowledge Representation Systems Based on Natural Language; (6) Learning Complex Behaviors in Adaptive Intelligent Systems; and (7) Plan Execution: Problems and Issues.

The American Association for Artificial Intelligence (AAAI) held its 1996 Fall Symposia Series on 9 to 11 November in Cambridge, Massachusetts. This article contains summaries of the seven symposia that were conducted: (1) Configuration; (2) Developing Assistive Technology for People with Disabilities; (3) Embodied Cognition and Action; (4) Flexible Computation: Results, Issues, and Opportunities; (5) Knowledge Representation Systems Based on Natural Language; (6) Learning Complex Behaviors in Adaptive Intelligent Systems; and (7) Plan Execution: Problems and Issues.

Configuration

Configuration combines parts into a satisfactory whole. The need for configuration is wide ranging—for computer systems to automobiles, telecommunication networks to insurance policies. A number of AI techniques have been brought to bear in this area, including case-based reasoning, constraint-based reasoning, description logic, genetic algorithms, and rule-based reasoning.

There was exceptionally strong industrial interest at the Configuration Symposium; over a dozen companies were represented. This meeting was the first of the configuration community on this scale, and discussion was lively. Discussion topics included (1) correctness, how we ensure that all and only correct solutions are produced; (2) process, how we capture and use restrictions on the process, for example, on the order in which decisions are made; (3) flexibility, how we accommodate different ways of specifying the problem; (4) maintenance, how we maintain knowledge in a configuration system over time.

Successes were identified: Real-world applications exist. Challenges were identified, especially in model acquisition and maintenance and in responses to overconstrained problems: providing explanations, suggesting alternatives, and computing optimal solutions.

Finally, some real community building occurred; a mailing list (configuration@lia.di.epfl.ch) and web site (http://www.cs.unh.edu/ccc/config/index.html) have already been set up, and plans are under way for a follow-up workshop.

Boi Faltings
Swiss Federal Institute of Technology

Eugene C. Freuder
University of New Hampshire

Developing Assistive Technology for People with Disabilities

Assistive technology is a growing field within the AI community. Some work in assistive technology can be viewed as an intermediate step to a full AI system; the inclusion of a human in the cognitive loop can allow solutions to be found for problems that have been unsolved until now. In addition, researchers in this area can see their research improving the lives of others.

The purpose of the Symposium on Developing Assistive Technology was to bring the community together for a weekend of discussion. This symposium grew out of a one-day workshop on assistive technology that was held at the Fourteenth International Joint Conference on Artificial Intelligence. Discussion centered on open problems, ways that the technologies should be moved into the target communities, and the importance of including the target community from the start of the research.

The keynote address was given by Paul Meyer of the President's Committee on the Employment of People with Disabilities. He stressed the need for community involvement from the beginning of the project; although many projects can technically be great, they will not be useful if they do not address a large segment of the target community. He addressed the needs of various communities and offered many suggestions on groups to contact for user involvement in the research process.

Symposium participants are conducting research in many different areas, including sign language translators, robotic wheelchairs, eye-tracking interfaces, assistive robots for the elderly and mobility impaired, and language aids. Many systems were demonstrated at the symposium. The papers were grouped into three main categories: (1) cognitive aids, (2) human-computer interaction aids, and (3) mobility aids. Participants gave a brief overview of their work and then presented three open problems in their research to open the discussion.

To foster more discussion, partici-

pants broke into work groups each day during lunch and then continued the discussion after brief presentations by each work group. On the first day, we discussed ways that we can move our research into the mainstream AI community and that we should continue to grow as a community. On the second day, the work groups focused on user interfaces. Most systems must be custom designed for each user. Participants were interested in finding a way to design our systems to allow for easy addition of these customizations. They also discussed how some individual systems could be put together for people with multiple disabilities.

The focus on discussions at the symposium allowed the participants to continue building their community. People exchanged resources for funding, user involvement, and collaborative research. A mailing list was created to continue discussions from the weekend. (Requests for additions to the mailing list should be sent to assistive-tech-request@ai.mit.edu.)

Holly Yanco
MIT Artificial Intelligence Laboratory

Embodied Cognition and Action

The Embodied Cognition and Action Symposium studied the role of embodiment and explored ways of extending existing low-level subcognitive systems such as autonomous robots and agents and grounding more abstract disembodied cognitive models. A large group of enthusiastic participants came from a broad range of backgrounds; the topics of paper sessions included philosophy, language and representation, motor control, neuroscience, robotics, and software-agent implementations. Each session started with a keynote speaker (Daniel Dennett, Jerry Feldman, Stephen Grossberg, Marcel Kinsbourne, Rodney Brooks, and Dana Ballard), was followed by short talks, and concluded with an open discussion.

Some heated arguments arose, but almost surprisingly, they converged on the last day, producing a set of agreed-on working definitions and principles: (1) Embodiment is a form of situatedness. (2) Embodiment versus situatedness is not the same issue as the real world versus simulation. (3) Embodiment stresses temporal and functional situatedness. (4) Embodiment enforces interaction with the agent's body and its environment. (5) Embodiment enforces energetic and resource considerations. (6) Embodiment prevents some inappropriate simplifications and unrealistic assumptions. (7) Embodiment enforces dealing with unexpected contingencies and provides specific dynamics. (8) Embodiment automatically grounds cognition (for example, provides natural biases for inductive models and representations). (9) Embodiment provides not only limiting, but also enabling, constraints for cognition. (10) Embodiment places the agent in the environment most similar to that of the designer and other natural systems. (11) Embodiment allows for direct cultural interaction. (12) Embodiment research should apply careful, selective pressure on the rest of AI. It should outline what was learned from working with embodied systems and how it affects the study of behavior and cognition.

Maja J. Mataric
Brandeis University

Flexible Computation: Results, Issues, and Opportunities

The Fall Symposium on Flexible Computation brought together an international mix of researchers studying issues with the elucidation and control of trade-offs in computational systems. Systems employing flexible computation have the ability to adapt their behavior to varying resources and challenges by modulating trade-offs between the quality of results and critical resources applied to problem solving, such as time and memory.

The symposium included a combination of panels, plenary sessions, and invited talks on principles and applications. Plenary sessions were structured from thematically related papers in the areas of planning and search, learning, new theoretical directions, and applications. Researchers shared their experiences with applications of flexible computation in search problems, digital signal processing, information retrieval, operating systems, graphics rendering, databases, face recognition and tracking, and robot path planning and execution. Attendees and invited guests participated in panel discussions on frontiers in utility-directed search, issues of representation, generalizable insights gleaned from diverse applications, and key open problems in flexible computation.

In the last sessions of the symposium, the attendees broke into focus groups where key problems were tackled. Groups explored challenging questions on opportunities for innovating with representation, the monitoring and control of computation, strategies for introducing new kinds of flexibility into traditional algorithms, tools and environments for research and development of flexible algorithms and systems, and the value of pursuing *bounded optimality*—optimizing the expected utility of performance given a set of explicit constraints on resources and architecture. Summaries of the results of the group discussions were shared in spirited presentations by group members. Attendees left the symposium enthusiastic about continuing to exchange results with their colleagues with common interests in flexible computation.

Eric Horvitz
Microsoft Research

Shlomo Zilberstein
University of Massachusetts

Knowledge Representation Systems Based on Natural Language

The Symposium on Knowledge Representation Systems addressed the theoretically and practically important problem of knowledge representation systems that closely parallel the representational and inferential characteristics of natural language. Natural language–based knowledge representation systems are attractive because they could automatically create and update knowledge bases from large corpora of texts in electronic

form and because such knowledge bases would be easy for people to use.

An interesting convergence of some views about the representation of meaning of natural language and knowledge representation systems (in theory and real-life natural language–processing systems) was reported. Many agree that natural language is like a knowledge representation system, a mental-level representation (a language of thought), and identify the representationally desirable characteristics of natural language. There is some consensus what these characteristics are, for example, a capability of representing and reasoning with ambiguous and disambiguated information.

A number of new and recently proposed inference methods motivated by natural language were discussed. One common theme in these methods is that a close correspondence of natural language syntax and structure of the inferable information results in natural language is both expressive and computationally tractable.

A number of researchers reported natural language–processing systems, existing and under development, that can learn and acquire knowledge automatically from a corpora of texts. These systems represent a new "natural language processing for knowledge representation" view of the relationship between knowledge representation and natural language processing, which complements the widely acknowledged "knowledge representation for natural language processing" view.

Lucja Iwanska
Wayne State University

Learning Complex Behaviors in Adaptive Intelligent Systems

Machine learning is an exciting area of intelligent systems research driven by scientific developments, interdisciplinary applications, and recent advances in computer technology that provide a unique opportunity to increase the role of learning in building complex knowledge representation systems.

The Symposium on Learning Complex Behaviors was a stimulating interdisciplinary meeting that brought together researchers from AI, adaptive control, neural networks, machine learning, qualitative modeling, language learning, and computational learning theory. Some of the highlights included plenary talks by Jim Albus and Les Valiant, who proposed different general architectures for complex systems, and panels that discussed new research directions for learning to model and control complex environments.

The participants then focused on new architectures, benchmarks, methodology, and algorithms for learning to perform complex behaviors and diverse cognitive tasks such as reasoning with uncertainty, planning, perception, language learning, and qualitative modeling.

The central theme of the symposium was research on agents that learn to behave "rationally" in complex environments, as typified by research on scaling up reinforcement learning using hierarchical methods or temporal abstraction (as proposed by Richard Sutton). Several talks stressed that traditional approaches to intelligent systems that separate the study of learning and reasoning (therefore yielding two computationally intractable problems) can be improved by a framework where an agent "learns to reason" relatively efficiently in a restricted context.

The symposium also addressed the need to create benchmarks that can be used for scientifically meaningful tests of the scalability and value of learning algorithms in complex systems and the issue of automation of *learning engineering*, that is, standardizing and minimizing user involvement in building adaptive systems.

Overall, this meeting was productive, addressing the scientific and engineering challenges facing researchers working toward increasing the role of learning in building knowledge representation systems that perform complex cognitive tasks.

Simon Kasif
Johns Hopkins University

Plan Execution: Problems and Issues

Just over 40 people participated in the Fall Symposium on Plan Execution. Most people were primarily interested in designing and building systems to make things happen rather than in studying the nature of cognition: Application areas included microwave heating of composites, telescope scheduling, and rocket control as well as the usual robots, both real and simulated. The mix of people made for many lively and wide-ranging discussions.

On the first day, there were sessions on the uses of formalism, adding deliberation into reactive systems and adding reaction in deliberative systems. A theme that emerged here was the contrast in emphasis between work directed at building practical applications, which tends to rely heavily on complex, expressive plan-representation languages, and work directed at integrating planning and execution, which use simpler plan representations that are easier to reason about.

The remaining sessions were on execution and scheduling, the exploitation of knowledge during execution, the matching of domain characteristics and architectures, and multiagent systems. It became apparent that many of the expressive activity descriptions that we had talked about incorporate a lot of domain and task knowledge. A major difficulty in using these task descriptions is the conversion of low-level sensor input into meaningful descriptions of the world. This conversion is usually so slow that responsive behavior is impossible. This problem is, of course, well known, and there are two common approaches: (1) use minimally interpreted sensor readings and (2) include explicit interpretation activities in the plans.

Louise Pryor
Harlequin Ltd.