

The 1997 Fall Symposia Reports

Sponsored by AAAI

Communicative Action in Humans and Machines

This symposium reexamined the view (proposed by Austin and developed by Searle and others) of communication as action rather than transmission of information. Such a view has become popular as a characterization of language use, and it plays a central role in the dialogue-management components of many systems that communicate with human users or other agents. An abstract level of representation such as speech acts is also useful as a media-independent characterization of the function of communication.

Current work that was presented and discussed at the symposium included both extensions to classical speech-act theory as well as attempts at standardization of speech-act labels. The extensions included accounts of dialogue phenomena other than classical illocutionary acts, such as turn taking, feedback, problem solving, and persuasion as well as the importance of social phenomena such as rights, roles, and obligations. Standardization groups are working toward the purposes of both inter-agent communication (using languages such as KQML) and the creation of sharable corpora of annotated dialogues (mostly between human dialogue participants) to allow language modeling at the dialogue level. An important topic is the evaluation of the reliability of such coding efforts across different coders, dialogues, and domains. The use of a proposed general-level coding scheme was used by the participants in a dialogue-coding exercise.

In addition, there were a number of presentations of human-computer

interaction systems engaging in communicative action using a variety of media. There was some interesting discussion about how the resources and constraints of machine communicative partners differ from human-human communication. Another important issue was the role of context, which motivated a joint session with the symposium on Context in Knowledge Representation and Natural Language Processing.

For more information (and links to some of the papers), consult the symposium web page: www.cs.umd.edu/users/traum/CA/.

David Traum
University of Maryland

■ The American Association for Artificial Intelligence held its 1997 Fall Symposia Series on 7 to 9 November in Cambridge, Massachusetts. This article contains summaries of the six symposia that were conducted: (1) Communicative Action in Humans and Machines, (2) Context in Knowledge Representation and Natural Language, (3) Intelligent Tutoring System Authoring Tools, (4) Model-Directed Autonomous Systems, (5) Reasoning with Diagrammatic Representations II, and (6) Socially Intelligent Agents.

Context in Knowledge Representation and Natural Language

We discussed current approaches to handling context necessary to simulate the apparently critical role of context in human knowledge representation, natural language processing

(NLP), and inferencing. Acknowledging huge differences in research goals and standards of knowledge representation and NLP, a diverse group of AI researchers (knowledge representation, NLP, linguistics, and philosophy) attempted to identify commonalities and differences.

Virtually everyone presented preliminary results; these were purely theoretical, with no (serious) implementation that would demonstrate the computational performance with respect to the theory (algorithm). Although some researchers never intend to implement anything, others quote serious difficulties in doing so, particularly on non-toy problems.

A common theme in the fields of knowledge representation and natural language is the idea that context acts as a filler of missing parts and that it simplifies and speeds up inference, thus compromising precision and correctness. In both fields, one observes a huge spectrum of answers to an important question in the technical agenda: "What is context?" This broad range of answers reflects both the confusion about context and the enormous difficulties in handling it.

The NLP-for-knowledge representation panel addressed the question of what knowledge representation-motivated research and development on context contributes to NLP and what the NLP-motivated research contributes to knowledge representation. It was widely felt that none of the existing knowledge representationlike theories of context handle or even address context-dependent aspects of natural language, such as semantic ambiguity, pragmatic inference of implicature, discourse structure, and underspecificity of negation. Little was said about the natural language-like approaches contributing to knowledge representation.

A simpler form of cross-fertilization of ideas was identified. A number of researchers pursue the application of existing theories from disciplines other than their own.

Lucja Iwanska
Wayne State University

Intelligent Tutoring System Authoring Tools

Intelligent tutoring systems (ITSs) are smart computer tutors and typically have an expert model, student model, instructional module, and interface. Although we could not always agree on the definition of these terms, we tended to agree that the separation of these modules is not always distinct, which becomes apparent in developing and using authoring tools for ITSs. In our symposium, we had attendees from 7 countries, and we saw 10 different authoring tools: (1) DIAG, (2) RIDES-VIVIDS, (3) XAIDA, (4) REDEEM, (5) EON, (6) INTELLIGENT TUTOR, (7) D3 TRAINER, (8) CALAT, (9) INTERBOOK, and (10) PERSUADE. In addition to demonstrating the tools, we compared and contrasted the components of the systems in discussions.

Some tools were meant for select authors or students. Others were designed for a wide set of authors. Some tools were designed to work with a limited area of domain expertise, and some were designed for a wide range of domains. Some tools had one main instructional strategy, but others had many.

Each tool had their own way of representing the student's knowledge and understanding of the material being taught. Some tools generated instruction directly from domain knowledge. Some relied on pedagogical knowledge about the domain to create instruction. Some provided simulation environments for practice and exploration.

We were struck by the incredible progress that has been made in ITS authoring tools and surprised that more tools and applications are not yet commercialized. From the symposium, we might develop a web site for ITS authoring tools.

We valued our time together, and we hope to keep having gatherings where we show and share our systems and ideas—maybe at ITS '98!

Carol Luckhardt Redfield
Mei Technology Corporation

Model-Directed Autonomous Systems

The information-gathering capabilities of the internet and smaller networked computational systems are offering new test beds for embedded autonomous agents. Physically, these agents involve a large distributed array of simple sensors, actuators, and processors. Functionally, their attention is directed inward toward maintaining their internal structure, although like traditional robots, they might also attend to exploring and manipulating their external environment. Controlling such systems is made difficult by the need to reason through a complex set of systemwide interactions; the one-of-a-kind nature of the test beds; and the need to coordinate a broad range of discrete, continuous, and software behaviors. These difficulties are being addressed by a new family of agent architectures, called *model-directed autonomous systems*, that use a compositional, declarative model to achieve the desired functions.

The symposium brought together a diverse set of researchers interested in exploring the concepts of model-directed autonomous systems. Participants within the symposium represented two different perspectives on embedded systems that rarely overlap: The first is the use of data-driven adaptive methods based on Markov decision processes and nonlinear control theory. The second is the use of highly deductive, symbolic methods for embedded control, developed within the model-based reasoning, planning, scheduling, and reactive execution communities. The symposium offered a unique opportunity for these two communities to develop a deeper understanding of each other's methods and to make progress toward a shared understanding of common research themes.

The symposium started with a discussion of key application areas for such systems and included presentations on life-support systems, chemical refineries, "smart matter," building energy systems, and deep-space missions. This discussion set the stage, identifying the key problems that drive the respective research communities.

This followed with a series of technical presentations on specific techniques related to a set of theme areas. One core theme that emerged was the use of different types of transition system model, including Markov decision processes, concurrent transition systems, and finite automata. These models are being used within embedded systems to provide a full range of functions, including planning, execution, fault diagnosis and recovery, and control. A related issue was the importance of using explicit probabilistic representations to cope with uncertainty. A second theme was the use of hybrid system models for control and simulation that merge symbolic deductive techniques with adaptive data-driven methods. These include the use of phase portrait analysis and the integration of discrete constraint-satisfaction techniques with continuous optimization techniques from operations research. The final theme related to the control of cooperative and adversarial multiagent systems and included the use of reinforcement learning to acquire system models.

P. Pandurang Nayak
RIACS, NASA Ames Research Center
Brian C. Williams
NASA Ames Research Center

Reasoning with Diagrammatic Representations

Diagrammatic representations can be defined as those that analogically model the semantics of a problem domain and *diagrammatic reasoning* as the process by which inferences are drawn from such representations. Since Larkin and Simon's (1987) seminal paper showed that explicit representation of a problem's spatial characteristics can result in an increase in computational efficiency, interest in diagrammatic reasoning, with a history stretching back to the beginnings of AI itself (for example, Gelertner's GEOMETRY MACHINE in 1959 or Evans's ANALOGY in 1962), has been on the rise.

Five years ago, this symposium was presented as part of the American Association for Artificial Intelligence 1992 Spring Symposium Series. This

symposium met with great success, generating over 100 high-quality submissions from a wide spectrum of disciplines, including psychology, philosophy, cognitive science, and AI. In its wake, a number of books and articles concerned with diagrammatic reasoning have been published (for example, *Diagrammatic Reasoning: Cognitive and Computational Perspectives*, Glasgow, Narayanan, and Chandrasekaran, MIT Press), a diagrams mailing list (diagrams@csli.stanford.edu) and web site (uhavax.hartford.edu/diagrams) have been initiated, and many research sites have turned their attention toward diagrammatic reasoning issues.

The current symposium was convened with the intent of consolidating research efforts since the original meeting and providing a forum in which to disseminate recent results and initiate new research. Keynote speakers B. Chandrasekaran and B. Tversky presented their perspectives—AI and cognitive science, respectively—on the field to date. Papers and posters were presented across a wide spectrum of disciplines, including those previously represented with the addition of visual programming, human-computer interaction (HCI), and information presentation. Two panels organized by R. Lindsay (DR Future Directions) and H. Narayanan (diagrammatic representations and HCI) were also presented.

Progress was made toward understanding the nature of diagrams and diagrammatic reasoning in that all attempts at definition were deemed inadequate. If a consensus was reached, it was that the notion of a diagram is far more subtle and complex than one would at first suspect and that the process of reasoning with them reflects this complexity.

Michael Anderson
University of Hartford

Socially Intelligent Agents

The symposium discussed sociality in software, robotic, and animal agents. Bringing together researchers from different fields resulted in cross-disciplinary discussions on how a single agent is embedded in a social and cul-

tural environment, how this agent interacts and communicates with other agents, and how societies of agents are formed. The majority of approaches referred to human-style forms of social interaction, which are, for example, required in agents that assist, cooperate with, or represent a human being. Despite technical and methodological differences in dealing with robotic and software agents, the symposium identified themes that cross the natural boundaries of agent species, for example, believability, narration, imitation, emotions, personality, cultural adaptation, and the coupling of internal and external dynamics.

A particular focus was on the role of the "human in the loop" as observer, designer, or user of social agents, for example, as a programmer of agent products, experimenter in robotics, and social interaction partner in software games and service robotics. Generally, agency and sociality are conceived of as characteristics of a system that can objectively be described and engineered. Current developments in areas such as believable agents, interactive art, personal software assistants, and virtual pets question this assumption and point toward works that have been done in philosophy, arts, cultural theory, and social sciences.

The symposium discussed both rational and irrational (emotional, subjective, inconsistent) aspects of socially intelligent agents, in this way stressing the need of a symbiosis between engineering and the humanities to build expressive, interactive, and social agents. This first Socially Intelligent Agents meeting did not attempt to achieve definitions of agent, sociality, and social agents. However, we started to discuss the complexity of social agents and appropriate design criteria in different applications. The symposium also addressed risks and opportunities provided by social agent technology.

Kerstin Dautenhahn
University of Reading

Expertise in Context: Human and Machine

Edited by Paul J. Feltovich,
Kenneth M. Ford, and
Robert R. Hoffman

Computerized "expert systems" are among the best-known applications of artificial intelligence. But what is expertise? The nature of knowledge and expertise, and their relation to context, is the focus of active discussion—even controversy—among psychologists, philosophers, computer scientists, and other cognitive scientists. The questions reach to the very foundations of cognitive theory—with new perspectives contributed by the social sciences. These debates about the status and nature of expert knowledge are of interest to, and informed by, the AI community—with new perspectives contributed by "constructivists" and "situationalists." The 23 essays in this volume discuss the essential nature of expert knowledge, as well as such questions as how "expertise" differs from mere "knowledge," the relation between the individual and group processes involved in knowledge in general and expertise in particular, the social and other contexts of expertise, the assessment of expertise, and the relation between human and computer expertise.

Published by the AAAI Press /
The MIT Press

590 pp., \$40.00
ISBN 0-262-56110-3

To order call 800-356-0343
(U.S. and Canada) or
617-625-8569.

Distributed by The MIT Press,
5 Cambridge Center,
Cambridge, MA 02142

Prices higher outside the U.S. and subject to
change without notice.