

AAAI 1992 Spring Symposium Series Reports

■ *The American Association for Artificial Intelligence held its 1992 Spring Symposium Series on March 25–27 at Stanford University, Stanford, California. This article contains summaries of the nine symposia that were conducted.*

Artificial Intelligence in Medicine

Unlike previous Artificial Intelligence in Medicine (AIM) meetings, the AIM Symposium was organized around a single source of medical knowledge, the 1991 article "Graft-Versus-Host Disease (GVHD)" by Dr. James L. M. Ferrara and Dr. H. Joachim Deeg (*New England Journal of Medicine* 324:667). Presenters were also given three clinical vignettes in GVHD to use as examples in their talks or poster presentations. The symposium was significantly more focused because of the use of a common knowledge source. Provocateurs had the task of asking difficult questions to both the presenters and the audience. Also, because most participants had carefully examined the article from the perspective of their own research, the depth of questions and discussion was vigorous and extensive. The symposium was dominated by two themes: knowledge sharing and reuse and temporal reasoning. Additional submissions on diagnostic reasoning and knowledge representation methods were also represented.

The symposium began with an overview talk by Peter Szolovits entitled "AI in Medicine: Past and Future." In this talk, Szolovits reminded the group that AIM has continued to make steady progress on difficult issues such as multiple diagnoses, probabilistic reasoning, and (simple) temporal reasoning. He also noted that like other AI application disciplines, AIM has moved away from a focus on core AI issues such as planning and learning. Although it remains a rich area for

research by AI specialists, AIM must now also address a broad and difficult interdisciplinary mix of topics, including information systems, data analysis, man-machine interfaces, pathophysiology, and genetics.

Theme one of the symposium was knowledge acquisition and knowledge reuse. Talks by Mario Stefanelli (Universita di Pavia), Jeff Bradshaw (Boeing Computer Services), Samson Tu (Stanford University), Nasir Amra (The Ohio State University), and William Punch (Michigan State University) focused on various methodologies for acquiring and reusing GVHD knowledge. Ramesh Patil (USC/ Information Sciences Institute [USC/ISI]) presented the Knowledge Sharing Initiative, and Michael Kahn (Washington University) contrasted this work with the American Society for Testing Materials 31.15/Arden Syntax work. The presentations and discussion on this theme illustrated that a significant amount of progress has been achieved in building tools that allow for the rapid representation and evaluation of qualitative simulation and rule-based models in the AIM community.

Theme two was temporal reasoning. Tom Russ (USC/ISI), A. Mete Kabakcioglu (University of Miami), Yuval Shahar (Stanford), Issac Kohane (Harvard University), Ira Haimowitz (Massachusetts Institute of Technology [MIT]), Barbara Heller (University of Bielefeld), and Gregory Provan (University of Pennsylvania) described markedly different approaches to encoding and reasoning about the dynamic aspects of GVHD. Provocateur Jim Hunter (University of Aberdeen) closely critiqued the various approaches to temporal reasoning. Hunter also noticed that most of the temporal ontologies were simple point-based methods that did not encode temporal uncertainty (a point also made by Szolovits in his opening talk). He also noted that

much work in the more traditional AI fields was relevant to AIM problems but that this research did not appear to be used by AIM researchers.

The remaining talks were dominated by diagnostic and knowledge representation methodologies: criteria-table methodology by May Cheh (National Library of Medicine) and Kent Spackman (Oregon Health Sciences), qualitative models by Serdar Uckun (Vanderbilt University), and a hierarchical knowledge representation methodology by H. Mannebach (North Rhine-Westphalia Heart Center). Researchers also presented a wide variety of methodologies during the poster session: Steve Cousins (Washington University) on the use of query networks, Peter Hucklenbroich (GSF-MEDIS-Institut) on the KLINC knowledge representation model, Donna Hudson (University of California at San Francisco) on a hybrid expert system model, Yeona Jang (MIT) on a causal graph diagnostic model, M. Magues (Institut European de Telemedecine) on a telemedicine assistant, John Weiner (University of Southern California) on a concept structure for representing consensus, and Jeremy Wertheimer (MIT) on the representing and the reasoning about molecular physiology.

The symposium was successful because all the participants had invested a significant amount of work in understanding a specific medical disease. It was unnecessary for each speaker to introduce his(her) specific problem domain. The examples used in the talks and posters were substantive because the presenters could assume that the audience had a similar knowledge of the example disease. Because the content of the symposium was greatly improved by this structure, it is hoped that future AIM meetings will adopt a similar model, perhaps selecting other articles that address segments of the AIM community that were

unable to participate in the GVHD exercise.

Michael G. Kahn
Washington University

Cognitive Aspects of Knowledge Acquisition

The symposium on Cognitive Aspects of Knowledge Acquisition originated in a controversy at the 1989 Workshop on Knowledge Acquisition for Knowledge-Based Systems in Banff, Alberta, Canada. At this 1989 workshop, one group of researchers promoted research on debiasing expert judgments and transferring undistorted knowledge, and another expressed that knowledge acquisition was not an activity of expertise transfer but rather of modeling, in which the expert and knowledge engineer collaborated in developing a model of the expertise that might not have existed previously.

Bill Clancey played a major role in focusing the resulting discussion and, as banquet speaker that year, regaled us with a hilarious presentation of a situated action perspective that has lived on in the memories of those who attended far more vividly than any academic paper. For Clancey, knowledge was not only in the community involved with the expert but also in the physical artifacts and environment in which the expertise was situated. Some tools so mold our behavior and promote our expertise that they are seen as vessels for knowledge, at least to the extent that our minds are seen to be so.

In subsequent meetings, the issue has never been as dramatically focused; meanwhile, the knowledge-modeling perspective has become widely adopted, and terminologies reflecting an expertise-transfer perspective have quietly been dropped. The symposium was taken as an opportunity to explicitly raise the issues again and, in particular, ask about the significance of the social situations in which expertise is situated, both theoretically and in terms of its influence on the design of knowledge-acquisition methodologies and tools.

The first session was based on a presentation by Bill Clancey, who recapitulated and updated the viewpoint he presented in 1989. One can best capture the essence of this presentation through a sequence of

aphorisms: Practice cannot be reduced to theory. You can represent knowledge, but the map is not the territory. Human behavior can appear regular without the person realizing that patterns exist or why they exist. To follow a pattern is not to interpret a recipe or a pattern *thing*. A representation of what a person knows is just a *model* of his(her) knowledge. Knowledge cannot be inventoried. What social practice makes possible our sense of a fixed, objective reality? What work does this objective view accomplish in maintaining and coordinating social interaction? How does creating and commenting on models change behavior?

The preceding paragraph effectively summarizes the meeting! The remainder was structured to promote discussion, with each 90-minute session begun with three short presentations aimed at developing critical issues rather than presenting research results. Discussants ranged from those concerned with the philosophical and psychological foundations of knowledge processes to those concerned with developing and using tools in practical situations to develop systems.

Steve Fuller discussed what social psychologists call the *fundamental attribution error*, not realizing that human behavior might be more a function of one's situation than of whatever knowledge one brings to the situation. Norman Livergood described work on modeling the behavior of politicians from their speeches and other information about them and using this information to predict their behavior in new scenarios. This paper triggered an unexpected theme of social responsibility in knowledge-acquisition research that recurred throughout the remainder of the meeting.

Steve Recogzei reminded us time and again that the knowledge underlying effective behavior is not necessarily tidy and nicely scientific: Lies and nonsense play important roles in many aspects of expert human behavior. John Boose brought us down to earth again with a paper that parodied various theoretical perspectives. His presentation reminded us that the knowledge-acquisition community came together with the primary objective of building tools, which are still the most convincing evidence of success—in theory or practice. His aphorism was that tools come first, but if we also have theo-

ries to support them, we should not be ashamed.

It is difficult to draw an overall message from what was a highly stimulating and discursive discussion meeting. For the tool builders, it was one of supporting the participatory process of building models from heterogeneous sources; modeling how knowledge structures change and not just the current state; and modeling knowledge itself as a wasting resource with local utility, not as an improving approximation to Platonic absolutes. One might also say that the message was to go on and do what we are doing but realize it is not just messy pragmatism; what we are doing might indeed be grounded in respectable theory.

Brian R. Gaines
University of Calgary

Computational Considerations in Supporting Incremental Modification and Reuse

The ability to modify previously synthesized artifacts, such as plans, designs, and programs, to meet new specifications is valuable in many tasks. Such incremental modification can be motivated by several considerations, including respecting previous commitments, minimizing change, and avoiding repetition of computational effort. Consequently, incremental modification has emerged as an active topic of research in many areas of AI, including planning, scheduling, design, and software engineering. The primary objective of this symposium was to bring together researchers working on modification issues in these different areas and facilitate a sharing of techniques and a cross-fertilization of ideas.

The symposium was organized around five presentation-discussion sessions on the following topics: reuse in planning and problem solving, analogy, software reuse, execution-time revision of plans and schedules, and dependency management support for modification and reuse. To cope with the diversity of these perspectives, the presenters were asked to characterize their work along a previously agreed-on set of domain-independent dimensions of modification tasks and techniques. The presentation sessions were com-

plemented by two invited talks by Jaime Carbonell and Jack Mostow and three panels that brought together the cross-disciplinary perspectives on incremental modification.

Much of the discussion revolved around understanding the spectrum of modification tasks, the utility of various modification techniques, and the bottlenecks for scaling up the existing modification techniques. Although a comprehensive account of the discussion is beyond the scope of this report, the following paragraphs summarize some of the broader issues that came up.

A consistent theme was the distinction between modification and reuse in the context of other first-principle techniques (such as generative planning) and modification and reuse where no first-principle method is available. Much of the work presented in the planning and problem-solving sessions fell into the first category, but the presentations on cross-domain analogical transfer, theory formation, and, to a large extent, software reuse and design modification all fell into the second category. It was noted that modification in the first case can formally be characterized as augmenting the problem solver with the ability to retract previous decisions. In the second case, modification needs to be either interactive or based on heuristic modification rules. This work led to the discussion of the scalability and the applicability of the techniques being used in highly constrained domains of the first case to more complex weak-domain theory situations such as software reuse. The design of effective retrieval strategies, rationale capture mechanisms, and generalizable domain-centric content theories of modification was recognized as the main bottleneck in this endeavor.

Another active topic of discussion was the factors affecting the utility of incremental modification and reuse. Regularity in the problems encountered in the domain, as well as the flexibility and conservatism of the modification strategy, was seen to affect the utility of reuse in a domain. It was argued that episodic learning of the kind supported by reuse could be superseded by more generalized knowledge as the agent's domain model becomes stronger, and the sample size of the problems encountered increases. (Mostow captured this intuition best by dubbing

stored plans as "poor man's control knowledge"). The place of reuse in the space of other speedup learning techniques (such as EBL) was discussed at length, and convincing cases were made for integrating reuse with these other techniques.

A third topic that received considerable attention was the content of the rationale structures stored along with an artifact to facilitate its later reuse. Based on accession complexity and the functions provided, distinctions were made between the causal structure of the artifact, the trace of the original derivation of the artifact, and the external rationales for individual design decisions made during the generation of the artifact. Some of these elements, such as causal structure, can be derived from the causal and functional models of the domains; however, other elements, such as derivational structures, must be cached during generation, limiting the latter to those situations where a generative problem solver can be relied on to annotate the artifact with the decisions made during the derivation.

Although the class of modification problems being addressed by the participants was too diverse to forge a unified theory of modification, the symposium was nonetheless successful in fostering a better understanding of the diverse bundle of techniques being used to aid the incremental modification of artifacts in various domains.

Subbarao Kambhampati
Arizona State University

Knowledge Assimilation

The main objective of the Knowledge Assimilation Symposium was to bring together researchers in machine learning and related fields to discuss their work from a shared perspective, with the hope of generating new ideas for moving forward in a common direction. The theme of knowledge assimilation was intended to focus attention on how a system whose performance improves in an independently changing environment must be capable of both acquiring fresh information and increasing the effectiveness with which it uses the information it already possesses.

In recent years, there has been significant progress on the two relatively separate tasks of accelerating

problem solvers and inducing concepts from examples. Important new techniques have emerged in both areas, notably explanation-based reformulation and the theory of probably approximately correct induction. Both these topics were discussed at the symposium—in formal paper presentations and informal conversations—but the focus was on combining and transcending speedup and concept learning, as exemplified by the title of one of the papers presented: "Concept Learning from Inference Patterns to Improve Performance."

The symposium attracted 40 participants from the United States, Japan, France, and Australia. Eighteen papers were presented, and two panel discussions were held. The papers spanned a wide range, from mathematical investigations of the inherent complexity of abstract learning tasks to reports on software used successfully in real-world applications. Most papers did not address the topic of knowledge assimilation explicitly, but some common themes emerged nevertheless. One theme was the need for an intelligent system that aims to improve its overall effectiveness to possess declarative knowledge about its own architecture, such as which classes of goals it achieves reactively and which by deliberation. Another theme was theory revision, which has probably become the most studied problem in machine learning research in the 1990s. This task is to use examples that a knowledge base (that is, theory) classifies incorrectly with examples that it classifies correctly to refine the rules in the knowledge base so that future examples are classified correctly.

The two panel discussions elicited considerable agreement and some controversy. The topic of the first panel was test beds for machine learning. Participants described opportunities for investigating knowledge assimilation issues in molecular biology domains, using mobile robots and *softbots*, agents that interact with existing software environments. There was considerable argument over whether it was fair to call these software agents robots. If any consensus emerged, it was that hardware robots face a radically harder task in that their sensors and effectors are intrinsically unreliable. The second panel discussion concerned the relationship between

knowledge representation and machine learning research. A number of participants identified concept-language restrictions that make problem solving tractable as an important point of common interest.

Charles Elkan

University of California at San Diego

Practical Approaches to Scheduling and Planning

Scheduling is the business of allocating time and resources to a set of tasks; *planning* typically precedes scheduling and is the business of defining the set of tasks to be scheduled. The automation of scheduling and planning has been studied for many years, but the recent focus on getting practical prompted the organization of this particular symposium. At the outset, we decided to use the term practical to refer to systems that explicitly reason about metric time, that is, systems in which one can state, for example, that a particular action must be done by 5 P.M. and must follow another by no more than 20 minutes. Although the implementation of scheduling and planning systems that manage metric time is definitely a nontrivial task, it is our contention that such systems can be expected to make an enormous difference to the overall efficiency of government and industry.

Symposium participants came from both research and operations organizations. Research participants were usually concerned with a general architecture or approach for solving a class of problems. Participants who came from real-world operational environments were mostly interested in solving a particular problem. This mix provided us with an excellent opportunity to match problem-solving approaches with problem instances. A large number of interesting application domains were represented, including telescope scheduling, airline scheduling, and power distribution management. As organizers, we established an atmosphere of collegial discussion early on (we interrupted a great deal), and this approach definitely raised the level of audience participation.

One major area of discussion was predictive versus reactive systems. A number of usefully different approaches were identified. Some systems are clearly based on the assumption that computation is suf-

ficiently cheap that it makes sense to simply react to each change or error as it occurs. Other systems are based on the idea that if information regarding possible errors is available in advance, then this information can be used to avoid errors or improve the system's response time when an error does occur.

A topic of constant debate was the apparent tension between so-called iterative improvement and constructive schedulers. An *iterative* improvement scheduler always has a complete schedule available for analysis. It proceeds by transforming this particular schedule into another one, still possibly flawed, until some termination criterion is met. In contrast, a *constructive* scheduler works from an incomplete schedule and repeatedly makes decisions about tasks and their time commitments (possibly backtracking at dead ends). Some participants felt that these two approaches represent distinct alternatives, and others felt that they represent different points on a yet-undefined continuum. One of the major points of agreement was that we need a better set of benchmark problems to help sharpen our understanding of the strengths and weaknesses of each approach. Some benchmarks do exist, but techniques are getting sufficiently good (in some cases) that existing problems do not help us to distinguish between them.

One particularly contentious topic was schedule robustness. *Robustness* relates to the ability of a schedule to withstand variations in execution, or at least some people thought so. Others felt that robustness was really a property of the entire system, of which the planner or scheduler was merely a part. This latter sense of robustness is possibly closer to that found in the control theory community; in this case, the schedule is simply a specification of what to do, and it might not make sense to refer to a given schedule as robust or not. Instead, under this view, what matters is the eventual set of behaviors produced by the total system. There were as many opinions as there were participants (if not more), and we agreed that robustness represents an interesting topic for future research.

We feel that the 1992 Spring Symposium on Practical Approaches to Scheduling and Planning was extremely successful in bringing together people from various technical disciplines to discuss recent

advances in the field. In addition, a number of interesting new ideas were generated, and we look forward to finding out how symposium participants work some of these ideas into the next generation of practical scheduling and planning systems.

Mark Drummond

Sterling Software

NASA Ames Research Center

Producing Cooperative Explanations

Many AI systems attempt to provide explanations in response to user questions. An explanation is cooperative if it is intended to help achieve the user's goals and be as easy as possible to understand. This symposium brought together approximately 30 researchers from 6 different countries, all of whom were worried about the problem of building systems that could produce cooperative explanations.

The traditional AI textbook view of explanation is an expert system answering a why question by providing an English version of the rule it is currently considering. The symposium participants, however, were working with a wide variety of systems, only a few of which could be classified as traditional rule-based expert systems, and only a few of which were providing system rules as part of their explanation. Participants were constructing systems to assist in software maintenance, advise novice UNIX users, answer database queries, generate textbook explanations of Lisp, answer questions about how various circuits work, teach students the skills of scientific argument, and assist expert system builders.

Explanation is a key component of all these systems. In building their systems, participants found themselves grappling with a wide variety of problems, including combining text and graphics in an explanation, providing examples as part of an explanation, recognizing when an explanation failed., repairing failed explanations, selecting appropriate phrasing for an explanation, evaluating the quality of an explanation, and determining when to ask questions as part of an explanation.

Despite this diversity, a general consensus quickly emerged on several key principles. First, a cooperative explanation is often a multimedia explanation, involving much more than filling in textual templates. This

viewpoint led to discussions of how systems could decide which information should be presented in what form and how systems could best mix different media. Second, a cooperative explanation is often a gradual, interactive dialog, not a one-shot, take-it-or-leave-it process. This discussion led to examining how a system could understand user feedback and revise explanations appropriately and studying different ways the system could be an active participant in an explanatory dialog, such as by asking questions of the user or actively critiquing the user's knowledge.

This community of researchers is small but active. They are taking the first few steps along what's likely to be a long road toward building systems that produce explanations that their users truly feel are cooperative.

Alex Quilici

University of Hawaii at Manoa

Propositional Knowledge Representation

The premise of this symposium was, "The key to propositional knowledge representation is that propositions can be represented by terms in a formal representation language, and hence properties of propositions and beliefs about propositions can be represented. This facilitates the study of representation and reasoning about beliefs, nested beliefs, and other propositional attitudes such as desires, wants, hopes, and intentions. Several knowledge representation formalisms based on the above ideas have been designed, proposed, implemented, and applied to various AI modeling tasks. Some examples include Sowa's conceptual graphs, Shapiro's SNePS, Arbab's propositional surrogates, and Wilks's ViewGen. Although the motivations for each of these may appear distinct, they all have to address a common core of knowledge representation issues. The goals of this symposium were to encourage a free exchange of ideas among the various groups of researchers, to discuss their solutions to common problems, to compare the theoretical and practical significance of their approaches, and to explore the possibilities for closer cooperation in the future."

Nineteen papers were presented. They ranged from position papers about the importance of propositional knowledge representation, identification of issues, and presentation of

actual systems employing ideas of propositional knowledge representation to methodologies for comparison of different systems. A selected subset of revised versions of these papers will appear in a forthcoming issue of the *Journal of Experimental and Theoretical AI*.

Approximately 30 people attended the symposium. We found, perhaps not surprisingly, that we were a congenial group with a wide overlap of interests and opinions.

We were all interested in the representation and use of the sort of knowledge that is transmitted among people in natural language. We assume the centrality of what are commonly called assertions, beliefs, or propositions, as opposed to the centrality of objects, classes of objects, and hierarchies of such classes. We were in general agreement that propositions are the objects of belief rather than sentences. In addition, although we found it hard to define precisely what a proposition is (just as it is hard to define just what a game is), we generally agreed that propositions are abstract entities in the domain of everyday thought and conversation and that a proposition is not equivalent to a sentence or a set of possible worlds. We were in general agreement that common-sense beliefs are important and that it is important for knowledge representation researchers to take natural language seriously as the mirror of the mind. We found a common interest in the nature and representation of the domain of discourse of knowledge representation formalisms, especially beliefs of other agents and incorrect beliefs.

Much of the agreement among the participants seemed to stem from experiences gathered from implementing actual systems. The common problems faced by everybody at this level were a unifying factor despite seemingly different theoretical approaches.

Stuart C. Shapiro

Hans Chalupsky

Deepak Kumar

State University of New York at Buffalo

Selective Perception

The world contains a tremendous amount of information—much more than an agent can process in real time, much more than is relevant for achieving its goals. Perceptual sys-

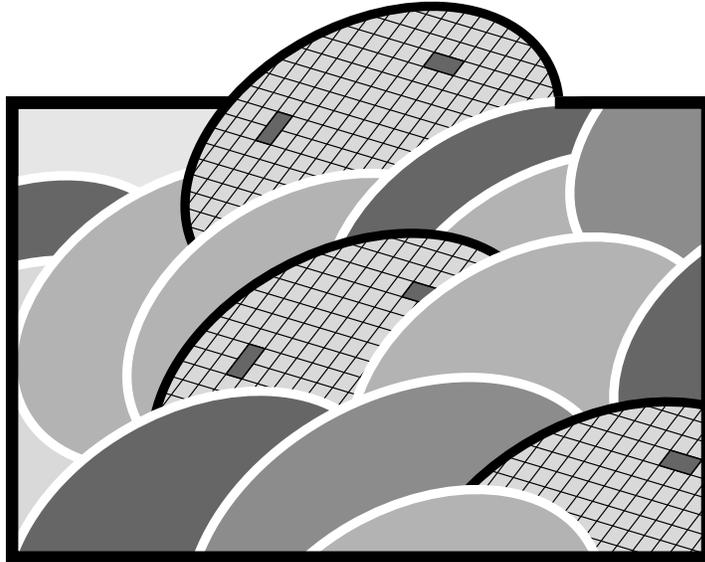
tems cannot hope to completely capture all the features in the environment, and planning systems cannot hope to be presented with complete and accurate descriptions of objects in the world. Research in selective perception is concerned with understanding what information is needed to perform given tasks and how perceptual systems can selectively be focused on certain aspects of the environment. Such task-directed guidance can include what features to sense, what resolution to sense at, where to sense, and what sensors to use for given tasks. The idea is that by focusing attention on the most relevant features, more efficient and accurate algorithms can be developed and used.

The symposium attracted researchers from a diverse set of fields, including robotics, machine vision, and AI planning. The opening speaker, Chris Brown from the University of Rochester, set the tone for the symposium by observing that the notion that perception is done for some purpose provides many constraints on the algorithms and system architectures that underlie intelligent systems. He pointed out that salience is a fundamental AI problem: whether it is understanding what is important for a robot to perceive or understanding what is relevant in a knowledge base.

The symposium format was based on panel discussions: The moderators prepared position papers that, together with rebuttals by panel members, were distributed to all the attendees. For each panel, the moderator presented the position paper, which was debated by three panelists (along with plenty of comments from the other attendees). Panel topics included the uses of decision theory, learning, planning, and biological analogues in selective perception research. It became clear that research in selective perception can benefit from the insights gained in these related fields. In particular, decision theory and learning provide focusing mechanisms for reducing large data sets to manageable size. A special session devoted to video presentations proved popular: Many impressive examples were displayed of robots and other intelligent agents selectively perceiving their environment.

Although general consensus was difficult to achieve with such a diverse group of researchers, one

AAAI-92 Conference T-Shirts Still Available



A few AAAI-92 conference T-Shirts (in the large and extra-large size only) are still available, postpaid for \$13.00, including tax. But hurry! supplies are limited. (Foreign orders, please add an additional \$5.00 postage. Call 415.328.3123 or send check to T-Shirts, AAAI, 445 Burgess Drive, Menlo Park, CA 94025.

INSTRUCTIONAL TECHNOLOGY FACULTY POSITION

Tenure track faculty position (half-time teaching and half-time research) is available. The teaching appointment is in the graduate program with special emphasis in interactive learning technologies. The research appointment is to participate as a senior scientist in association with Dr. M. David Merrill and the ID₂ group. This individual will be expected to generate funding for the research appointment. Some international travel required. Requires an earned doctorate (preferred areas of specialization include instructional technology, cognitive science, computer science, or a related field), significant experience in the application of computers to instruction, demonstrated research skills and a record of professional publications. Send letter of numbers of five references to: Dr. M. David Merrill, Department of Instructional Technology, Utah State University, Logan, UT 8432202830. Review will begin on September 1, 1991, and will continue until an acceptable candidate is selected. Utah State University is an Affirmative Action/Equal Opportunity Employer — Women and Minorities are strongly encouraged to apply.

common theme was the need to tailor perception to fit the requirements of the tasks. However, the problem of defining what a task is remains; this discussion led to debates on the utility and feasibility of formalizing tasks and their information requirements. The general consensus was that we are not far enough along to do so: Each new system adds new techniques, but overarching principles have not yet emerged. However, a major objective of the symposium—establishing dialogue among interested parties—was accomplished, leading us to expect exciting advances in the future.

Reid Simmons
Carnegie Mellon University

Reasoning With Diagrammatic Representations

The symposium on reasoning with diagrammatic representations was an interdisciplinary gathering that focussed on imagery, diagrammatic representations, and reasoning using these representations. It brought together researchers from disciplines as varied as psychology, logic, philosophy, and computer science. The main goals were to initiate cross-discipline dialogues, to identify and debate important issues, and to further fuel research interest in this area. A detailed report on this meeting is being written.

N. Hari Narayanan
Ohio State