



Guest Editor's Introduction

The 1990 Conference on AI Applications

Chris Tong, Rutgers University

Like its predecessors, the 1990 Conference on AI Applications was devoted to the application of artificial intelligence techniques to real-world problems. The program committee solicited two kinds of papers:

(1) Case studies of knowledge-based applications that solve significant problems and stimulate the development of useful techniques.

(2) Papers on AI techniques and principles that underlie knowledge-based systems and enable more ambitious real-world applications.

In the past, many papers submitted to CAIA had involved engineering and manufacturing. The 1990 program committee wanted to encourage a more balanced picture of AI applications, so the chair, Se June Hong, created three tracks: engineering/manufacturing, business/decision support, and enabling technology. The first two tracks highlighted specific applications. Industrial and scientific applications appeared in the engineering/manufacturing track, while applications from business, government, and law were slated for the business/decision support track. Authors in these two tracks were required to

- justify their use of AI techniques based on the problem definition and an analysis of the application's requirements,
- explain how they used AI technologies to solve a significant problem,

- describe their implementation's status, and
- evaluate both the effectiveness of their implementation and the technique used.

The third track focused on more application-independent tools, techniques, and principles, specifically those aiding the development of practical knowledge-based systems that can be scaled to handle increasing problem complexity.

The review process was more stringent for the 1990 conference than before. Of the 192 papers submitted, only 44 were accepted. Most papers received five reviews, and program committee members thoughtfully discussed any differences among reviewers. While most of the accepted papers were from the United States, research from other countries was represented as well: five papers were from Canada, four from Europe, and three from Japan. The acceptance rate for each track was as follows:

- Engineering/manufacturing: 23 of 83 papers accepted. Topics included diagnosis, design and planning, scheduling, and space applications.
- Business/decision support: Nine of 48 papers accepted. Topics included inductive generalization and domain-specific tools.
- Enabling technology: 12 of 61 papers accepted. Topics included domain-independent tools, user interfaces, and advanced rule systems.

There were also combined track sessions with topics including connectionism in AI, abstract models, office systems, and project management.

The type of papers at CAIA 90 differed from those at AAAI or IJCAI. Very little "formal" research was presented, perhaps because of the program committee's emphasis on verifying and understanding systems. Also, the presentations reflected greater synergy between research and applications, particularly those in the enabling technology track.

Several research topics were highly represented — such as diagnosis, planning and scheduling, and design — while others were notably absent. The lack of material on AI and software engineering, AI and deductive databases, and natural language processing (the exception being David Waltz's plenary address) may reflect either the availability of more specialized conferences for these topics or the fact that current applications do not use much work from these areas. Another explanation is that researchers in these areas are not aware of CAIA or do not feel it is an appropriate forum for their work. (The News department in this issue looks further at CAIA's future; see p. 52.)

Talks

One of the most common remarks at the conference was, "There are an awful lot of AI applications out there!" Mark Fox's remarks in the conference proceedings supported that impression: "Over 1,000 systems have been documented to be in production use, and it is estimated that over 3,000 systems are in production use today. (DuPont claims that they alone have about 600 systems in use.)"

Ed Mahler's relatively nontechnical keynote address — "Knowledge-Based Systems: The Competitive Imperative of the '90s" — had a message for everyone. For active researchers, it offered arguments to convince managers and funding agencies of the need to investigate knowledge-based systems. For nonresearchers, it provided reasons to believe in such systems.

The point of the plenary address by David Waltz,

The IEEE Expert special track

CAIA's stringent review process in 1990 ensured that the papers were generally of high quality. Of the accepted papers, Peter Hart and I selected four that were sufficiently representative of the best work in AI applications that they should be expanded into articles. Beginning in this issue of *IEEE Expert*, these articles will appear as part of a special track devoted to AI applications. (Alex Chao-Chiang Meng and Michael Sullivan's *Logos: A Constraint-Directed Reasoning Shell for Operations Management* appears in this issue, as does *Resource Allocation in Distributed Factory Scheduling* by Katia P. Sycara, Steven F. Roth, Norman Sadeh, and Mark S. Fox. Future issues will feature Steven K. Feiner and Kathleen R. McKeown's *Generating Coordinated Multimedia Explanations* and Hwee-Tou Ng's *Model-Based, Multiple-Fault Diagnosis of Time-Varying, Continuous Physical Devices*. Ng's paper won the Best Student Paper Award at the conference.

"Massively Parallel Text Processing," was that document retrieval can be done very efficiently on a massively parallel computer such as the Connection Machine. This project seems to be going well, and Waltz provided good performance figures. However, he failed to point out that, although retrieval speed increases as the number of processors increases, the speed of updating such a knowledge base does not. Waltz also has not done much work in natural language text processing, a much tougher area that needs more research. However, Waltz's company has

tested his ideas on Wall Street, and one program committee member suggested that such research is just the kind that should be reported in the business/decision support track at future conferences. Such work relates to real business problems and concerns that real executives face.

In his invited talk on "Modeling and Analyzing Business: A High-Payoff Challenge for AI," Richard Fikes spoke of the potential for exploiting qualitative reasoning in analyzing financial data. This is a promising area in which several attendees were optimistic that progress would be made.

Panels

There is a common impression that AI researchers develop prototype ideas and methods but leave to others the tasks of making them work in a larger context, making the algorithms efficient, and then efficiently implementing the algorithms. Thus, the panel on "How to Deploy Lisp Applications" was unique in that it showed AI people starting to think about deploying truly big applications. Many attendees were interested in the scale-up issues, and several ideas were tossed around, including ways to speed up Lisp, investigations of other programming environments such as C or C++ (though most attendees agreed that Lisp could be an appropriate vehicle with a little more effort), and a package for converting Lisp programs into Ada. The panel conveyed that a lot of work is being done on improving environments, development tools, compiler efficiency, etc.

Other panels questioned whether Lisp is a "serious,"

The papers differed from those at AAAI or IJCAI. Very little "formal" research was presented, and there was greater synergy between research and applications.

commercially viable AI language. In the plenary panel, "The Future of AI and AI Applications," Mark Fox declared Lisp "dead" and supported his view by pointing out the significant number of Lisp-based shells recently rewritten into C. On the other hand, the panel titled "How to Deploy Lisp Applications" made some positive points about the language's viability, including performance benchmarks comparing Lisp and C (ignoring memory images) favorably.

The panel on "New Funding Initiatives for AI" provided not only valuable information on funding, but also interesting technical ideas. Bob Kahn, president of the Corporation for National Research Initiatives and former official of the Defense Advanced Research Projects Agency, discussed a national digital library concept connecting many digital libraries over networks and including not only traditional elements of libraries, but also software, sensor data, and databases. The library would incorporate "knowbots" (agents to help users do library searches) and indexing by "icon geography." It would also have to deal with such issues as converting knowledge or data from its given form to the form needed by the user; understanding the connections between knowledge structures and printed formats; and surmounting diverse barriers (economic, technical, social, legal, standards, critical mass, business practices, and risk). In this last area, intellectual property issues would be the hardest. Kahn compared the envisioned library to existing large-scale enabling technologies like the national highway network, power generation and distribution networks, and the telephone network.

Daniel Masys of the National Institute of Health described the human genome project and its computing needs. The human genome comprises 3×10^9 base pairs. It now costs \$1 per base pair to sequence the genome, and current articles report around 10^4 base pairs sequenced. Improved technology will reduce the cost and allow the whole genome to be sequenced. What will then remain is the very large information-processing problem of predicting structure and function from the base pair sequence. (For example, molecular biologists know the DNA codes for amino acid sequences that form proteins, but not how the proteins will fold, which is key to their function.) About 10 to 15 percent of the funding is earmarked for informatics. Therefore, the way to get funding in this area is to have an interdisciplinary team with both molecular biologists and computer scientists. One attendee pointed out that if the genome sequence were available on CD-ROM, looking for interesting

patterns could become a cottage industry; the disk would be cheap, and workstations would have enough power to process it.

Bob Simpson of DARPA discussed the Federal High Performance Computing Program. This "meeting of minds" from eight federal agencies on the importance of computing covers such areas as high-performance computing systems, advanced software technology and algorithms (the so-called Grand Challenges), the proposed National Research and Education Network (NREN), and basic research and human resources. The program is not yet funded, but might be by fiscal year 1992. Proposed funding for the first five years is \$151 million, \$256 million, \$411 million, \$602 million, and \$597 million. [A special track on this program will appear soon in *IEEE Expert*. — Ed]

Acknowledgments

In this article I have relied greatly on observations by Kevin Ashley, Vasant Dhar, Bob Engelmores, Walter Hamscher, Se June Hong, Sitaram Lanka, Diane Litman, Eric Mays, Hwee-Tou Ng, Gordon Novak, Katia Sycara, the 1990 AI Application Conference Program Committee, and many people with whom I chatted at the conference.



Chris Tong is an assistant professor of computer science at Rutgers University. He was the program chair of the engineering/manufacturing track of CAIA 90 AI, and he is the program chair of the scientific/engineering track of CAIA 91. He has presented several tutorials on AI and design, including at the 1988 National Conference on Artificial Intelligence (AAAI 88) with Lou Steinberg, and at AAAI 90. He will present tutorials at AAAI 91 with D. Sriram, and at the 1991 AI Applications Conference (CAIA 91) with D. Navinchandra. He also organized a workshop (with D. Sriram) on AI and Engineering Design at the 1986 AAAI Conference.

Tong was guest editor for the July 1987 special issue on artificial intelligence in engineering design of the *International Journal for AI in Engineering*. He is a coeditor (with D. Sriram) of the forthcoming three-volume collection, *AI in Engineering Design*.

Tong earned his bachelor's degree summa cum laude from Columbia University in 1978, and his PhD from Stanford University in 1988. He was the recipient of the 1990 Rutgers CSGSS (Computer Science Graduate Student Society) Award for Excellence in Teaching.

Readers can reach Tong at the Department of Computer Science, Rutgers University, Busch Campus, New Brunswick, NJ 08903; e-mail: tong@cs.rutgers.edu