

The Innovative Applications of Artificial Intelligence Conference: Past and Future

Howard Shrobe

■ This article is a reflection on the goals and focus of the Innovative Applications of Artificial Intelligence (IAAI) Conference. The author begins with an historical review of the conference. He then goes on to discuss the role of the IAAI conference, including an examination of the relationship between AI scientific research and the application of AI technology. He concludes with a presentation of the new vision for the IAAI conference.

The eighth annual meeting of the Innovative Applications of Artificial Intelligence (IAAI) Conference was held on 5 to 7 August in Portland, Oregon. Over the past eight years, this conference has undergone modest evolution, but a significant transformation is being planned for the next meeting. Therefore, it seems an appropriate time to reflect on the goals and focus of the IAAI. I begin with an historical review of the conference in which I examine the content, the organizational forms, and the themes that have dominated the various meetings of the conference over its eight years. I then turn to the topic of what role the IAAI conference serves in the overall mission of the American Association for Artificial Intelligence (AAAI). This discussion involves examining the relationship between AI scientific research and the application of AI technology. I review how the past chairpeople of the IAAI have understood the mission of the conference and then suggest that an important part of the mission has not adequately been addressed. Finally, I conclude by presenting the new vision for the IAAI conference.

After many years of high ambitions and optimistic predictions, by the mid-1980s, it had become clear that AI research was leading to successful applications. The IAAI was

formed to highlight this successful transition of AI technology from theory to practice, recognize AI applications and AI application developers as integral contributions and contributors to the AI field at a national conference, and provide a forum for the exchange of experiences and lessons learned in the heartland of the AI community. It was desirable to show that AI technology was capable of delivering significant value to both commercial and government institutions, show the breadth of application areas, and show that the applications enabled by AI technology were revolutionary. The name of the conference was consciously chosen to highlight this revolutionary aspect of AI technology: An *innovative* application is one in which AI technology had enabled solutions for problems not previously thought to be amenable to computational solutions.

Table 1 shows a categorization of the applications presented over the history of the conferences. Three things are striking: First, these applications cluster in areas that are central to a broad segment of corporations. Second, the systems cover an enormous diversity of application areas. Third, applications peculiar to government institutions (military or space) represent a significant but nevertheless relatively minor subset. (However, many of the applications in other categories were delivered by governmental institutions).

Looking over the eight years of the IAAI, there seem to have been three periods, each characterized by a dominating theme. In years 1 and 2 of the conference, the dominating theme was "AI, it's for real." The introduction by Alain Rappaport and Reid Smith (1991, p. xi) to the second IAAI proceedings

Table 1. A Categorization of the Applications Presented over the History of the IAAI Conferences.

Manufacturing and design	30
Business operations	26
Finance	25
Telephony	12
Diagnostics and troubleshooting	11
Claims processing and auditing	10
Information retrieval and classification	9
Computers and software engineering	8
Military	6
Space	4

highlights this theme:

While academic research continues to progress, **deployed applications demonstrate the utility** of this new computer technology. Applications have become a driving force in **bringing AI into the mainstream**, well integrated with existing computing environments (*emphasis mine*).

Certainly, part of this theme was illustrated by the deployment of *elephant applications*, those that were large scale and produced dramatic payback in both relative and absolute terms. Among these were applications from such large organizations as IBM, AMEX, Digital Equipment, Chase Bank, and the United States Navy. The applications were not limited to the United States; several came from Japan, Singapore, and Korea. Not all the applications were elephants; there were many *rabbits*, small-scale applications of AI, built quickly and nimbly, with modest investment and significant relative (but not necessarily large absolute) payback. Dupont, in particular, championed the strategy of building many rabbits rather than a single large elephant.

The greatest impediment to getting these applications built was the difficulty of integrating the AI component of the application with the conventional flow of business data processing. It was this integration task that became the dominating theme for the next several years (years 3–6 roughly). As program chairs Reid Smith and Carlisle Scott (1992) noted in the introduction to the proceedings of IAAI 3,

On a **discouraging note**, application developers continue to be forced to spend large fractions of their time **cop-**

ing with unfortunate and grim realities of integrated computing in today's world.... For too much **talent is dissipated** in piecing together, rewriting or **force fitting low-level components** not designed for evolution (*emphasis mine*).

What caused the problem was a mismatch between the platforms that were used to develop AI applications and the platforms that supported mainstream business data processing. The AI components were typically built to run on UNIX workstations (and Lisp Machines in the earlier days) that were comfortable in a distributed, internetworking environment. In contrast, the corporate mainframe environment was more proprietary, provided inelegant communications and weak support for distributed processing, but it owned the data. The usual approach to integration was to make the workstation pretend to be a 3270 dumb terminal; it would receive the same textual forms that the conventional application developers had intended for a human operator and then extract the information useful for the AI application from these textual forms, a technique known as *screen scraping*.

Because AI was a leading-edge technology, it arrived in this world too early. As a consequence, the AI application community had to ride many waves of technological quick fixes and fads. Over the years, the focus moved from corporate mainframes with their Cobol and PL/1 orientation to PCs, C, and C++; proprietary client-server frameworks dominated for a while. Unfortunately, keeping up with these changes seemed to occupy more of the application developers' time than keeping up with new research results from the AI research community.

Many of these integration problems are now being addressed head on by a broad community of information technologists using Internet-based frameworks such as CORBA and the World Wide Web. Ironically, the network orientation of the early workstation-based systems (which earlier seemed like an impediment) is exactly what everybody is now scrambling to regain. I do not mean to argue that the integration problem was unimportant; clearly, it was critical. However, it is a transient and ever-changing problem. It would be an exciting prospect to see AI technology applied to this integration problem, so that what is now a painful and low-level-system programming task might, in the future, be reduced to the much simpler task of describing the characteristics of the systems that need to be integrated. I look forward to

seeing such intelligent integration technology reported on in future IAAI conferences.


In the last couple of years, as tools for integrating with enterprise-scale computing have become more common in the mainstream, and the technology for building AI systems has become relatively routine and mature, the focus of the IAAI meetings has again shifted. Now, the case studies in the conference focus primarily on the domain and task analysis and on the appropriate matching of the problem characteristics to available technological solutions. I find this a refreshing trend and hope that it truly characterizes the spirit of upcoming IAAIs.

The organizational form of the IAAI conferences mirrored to some extent these themes. The first two conferences were held in separate locations and different times from the national conference on AI to call special attention to the exciting new applications. Over the next several years, the IAAI conference was moved into greater alignment with the national conference. First, the two conferences were colocated but maintained separate registrations, proceedings, and reviewing policies; in later years, certain sessions were shared between the conferences, but in general, the feeling was that the IAAI and the national conference should maintain separate identities. Because the dominating issue for IAAI papers was often integration with mainstream computing, there was limited opportunity for interaction with the AI research community.

As the focus of the IAAI has returned in the last couple of years to domain and problem analysis, it has seemed increasingly relevant to integrate the IAAI more completely with the AAAI conference. In 1996, we moved to a single registration fee and a single proceedings, including the papers from both conferences. Next year, as the call for papers makes clear, we move even further in this process of integrating the two conferences.

Given that the IAAI and AAAI conferences are moving closer both intellectually and organizationally, it seems appropriate to turn now to a discussion of what the nature of AI research is and how it relates to the practice of building applications of AI technology.

As a scientific discipline, AI has always been a bit unusual. According to at least one definition, AI is the study of machine intelligence. Under this definition, AI is quite unlike the natural sciences, which study existing phenomena; AI studies the (largely) nonexistent phenomenon of machine intelligence and seeks to bring it into existence. AI has



always, therefore, been both speculative and constructive: We don't necessarily know what a machine intelligence would look like, so we try to build one, perhaps learning more from our mistakes than from our successes. This experimental approach to AI science is hard to distinguish from the more pragmatic engineering attempt to build useful programs that exhibit some degree of intelligent processing.

Of course, AI is not just the attempt to understand machine intelligence; for many, it is equally concerned with understanding human intelligence, either as an end in itself or because human intelligence provides the most useful model of intelligent behavior that we can find. In this view of AI science, we build programs as experimental models of human cognition. In building such models, we often find ourselves asking what it is in the way people think about a certain problem that allows them to effectively cope with it. Again, these scientific pursuits are hard to separate from the similar engineering pursuits of building programs that do useful tasks by emulating human cognition. In the scientific pursuit, we ask how people look at a task in

As the focus of the IAAI has returned in the last couple of years to domain and problem analysis, it has seemed increasingly relevant to integrate the IAAI more completely with the AAAI conference.

IAAI Articles in This Issue

This issue of *AI Magazine* includes five case studies that were presented at the Eighth Innovative Applications of Artificial Intelligence on 5 to 7 August 1996 in Portland, Oregon.

“Diagnosing Delivery Problems in the White House Information-Distribution System” by Mark Nahabedian and myself describes an AI component of a much larger system that the AI Lab at the Massachusetts Institute of Technology has built for use by the White House Office of Media Affairs. When the Clinton administration took office, it decided to utilize the rapidly growing National Information Infrastructure to bring information directly to the average citizen. Speeches, press releases, and the transcripts of press briefings are all made available daily to anybody interested through both the World Wide Web and e-mail. The e-mail channel allows users to create a profile specifying the type of information that interests them; as documents are released, they are e-mailed to all those whose profiles indicate an interest. This system has been estimated to reach more than 100,000 recipients. Unfortunately, the Internet is a relatively volatile and anarchic environment: E-mail addresses become invalid frequently as users move from one site to another. The bounced mail expert system (BMES) described in this article diagnoses such failures to get e-mail delivered and helps the administrator remove obsolete mail addresses. The system has reduced the portion of the system administrator’s daily work load consumed by the management of bounced e-mail from 3 hours to 30 minutes.

“Intelligent Retail Logistics Scheduling” by John Rowe, Keith Jewers, Joe Sivayogan, Andrew Codd, and Andrew Alcock describes a scheduling system developed by J. Sainsbury, one of England’s largest retailers (with 11.7 percent of the market and 12 billion British pounds of annual business). The supply-chain integrated ordering network (SCION) described in this case study has allowed Sainsbury to dramatically increase the efficiency of its depots and distribution network for nonperishable goods. The system manages 22 warehouses and processes between 100,000 and 200,000 pallets of nonperishable commodities to be placed on 5,000 to 10,000 vehicles each order cycle. SCION is, in effect, a constraint-directed scheduling system that determines, among other things, when goods should be delivered and which vehicles should carry them. SCION has allowed Sainsbury to move from weekly planning and scheduling to daily; Sainsbury estimates that the system will return its investment in 6 months and will lead to benefits of more than 10 million British pounds over the next 5 years.

“Using Artificial Neural Networks to Predict the Quality and Performance of Oil-Field Cements” by P. V. Coveney, T. L. Hughes, and P. Fletcher describes an innovative use of neural network technology in oil-field operations. Schlumberger is one of the world’s largest oil-field service companies, and it was also one of the first large companies to explore the use of AI technology. Oil-field cements are used to line oil and gas wells after drilling by pumping a slurry between the well bore and a steel case inserted into the well. The thickening time of the cement is critical: If the cement sets too quickly, it can set in the pumping equipment or otherwise cause catastrophic failure; if it sets too slowly, then it imposes an unnecessary delay until drilling can resume. Delays can cost more than \$1 million a day. Unfortunately, the properties of cement are extremely difficult to predict. Schlumberger tackled this problem by using a

such a way that makes it computationally tractable, and then we go on to ask what abstractions, knowledge, problem-solving methods, and computational machinery support this. In the engineering task, we ask how people look at the task we’re interested in automating, and then we go on to examine what abstractions, knowledge, problem-solving methods, and computational machinery might support it. We then go on to make pragmatic decisions about both what we know how to automate economically and what is institutionally relevant to automate.

The science of AI and the engineering discipline of AI have much to learn from one another, and it is this synergy that we hope to further develop with the new format for the IAAI. To date, we have seen one small example of how the conference can facilitate interchange between the scientific and engineering aspects of AI. Case-based reasoning was a new problem-solving paradigm that emerged from the research environment during the mid-years of the IAAI. It was one picked up quickly by the application development community because of its relevance to help desks and a large number of other applications. However, there is a significant gap between the practical uses of case-based reasoning, which are limited almost totally to case retrieval, and the full paradigm as it emerged in the research community, which involves not just retrieval but also case adaptation, debugging, and so on. The research antecedents of this paradigm go far back in the AI research literature to early work on planning (STRIPS) and debugging (Sussman’s HACKER) as well as to the more obvious antecedents in the work of Roger Schank and his colleagues on scripts and other episodic representations of knowledge.

Why didn’t this rich research tradition show up in the applications? There are several explanations, not necessarily contradictory: First, in many application settings, case adaptation is what people do, and they’re good at it; retrieval is what machines do better. A second explanation is that the application areas pursued don’t require case adaptation. A third explanation is that the research in this area isn’t yet ready for prime time, and we need to better understand the paradigm, particularly in the context of demanding applications. The fact that this third explanation is at least plausible led to a AAAI fall symposium on knowledge adaptation, perhaps with some useful outcomes.

How then did the organizers of the IAAI conferences conceive of the goal of the con-

ference? We reviewed the program chairs' introductions to the proceedings of the IAAI conferences to get a sense of this. We found the following comments in the introduction to the first proceedings (Schorr and Rappaport 1990, p. vii); notice its emphasis on the "it's for real" theme:

The **emergence of scientific achievements** from the field of AI has triggered **opportunities to tackle new problems** in which the computer formerly had no role.... If we have not achieved all the **goals of automating cognition yet**, our years of research are rewarded by what we feel can have a **major impact on everyday operations in the workplace...** (*emphasis mine*).

Program chairs Reid Smith and Carlisle Scott (1992, p. ix) take this theme even further in the third conference proceedings, describing the conference as "a casebook for the **businessperson**—showing in a variety of domains **what can be done with the current technology** as well as **practical problems** that arise in its application" (*emphasis mine*). I already noted previously that the introduction to the third conference emphasized the integration problem (bemoaning its "grim realities"), which dominated the conference for many years.

However, we also find in these introductions inspiring ideas about how the scientific and engineering aspects of AI can relate to one another. The first conference proceedings (Schorr and Rappaport 1990, p. vii) notes that "[s]uch a forum can further **the exchange of information on what really works and what the problems are**. Tackling these new applications will **lead to better technology** because we will **find and remedy current deficiencies** when solving 'real' problems" (*emphasis mine*). The introduction to the third conference (Smith and Scott 1992, p. ix) suggests that the research community can find its best challenges in the applications world: "**Research is driven by applications**—especially those pressing on the frontiers of what we know how to design and implement" (*emphasis mine*).

What then should be the role of the IAAI? Conference organizers believe that the correct emphasis going forward is in the vision represented in these two quotations, that is, to "further the exchange of information on what really works and what the problems are" (Schorr and Rappaport 1990) with the hope that "tackling these new applications will lead to better technology" (Schorr and Rappaport 1990). AI research and the applica-

combination of powerful AI and signal-processing techniques: They trained a neural network to use a Fourier transform infrared spectrogram of the cement to make accurate predictions of its properties. This combination of technologies has not only led to greater accuracy, but it has also contributed to the understanding of cement chemistry.

"COMET: An Application of Model-Based Reasoning to Accounting Systems" by Robert Nado, Melanie Chams, Jeff Delisio, and Walter Hamscher describes an advanced system being deployed at Price Waterhouse, one of the world's largest accounting firms. One of the services that Price Waterhouse's auditors provides its customers is an analysis of the controls built into the customer's accounting systems. Accounting controls attempt to guarantee that the data within the system are accurate and nonfraudulent. Accounting systems consist of both automated and manual processes; in today's world, they are large, distributed information systems whose design is not often available. Thus, one of the auditors' most difficult tasks is to construct a model of the accounting system that is accurate enough to assess the efficacy of the controls within the system. COMET provides both model-building and model-analysis support. The construction tools help by providing a convenient, graphic, and hierarchical framework for developing the models and consistency-checking tools to catch obvious errors. The analysis tools are based on the observation that all actions in an accounting system can be reduced to one of a few primitive action types and that each of these is subject to only a few basic failure types. COMET can simulate these primitive failure types and then use the model to determine if these failures will be caught by controls within the accounting system. To date, more than 20 percent of the Price Waterhouse auditors have been trained in the use of COMET, and it is rapidly becoming the recommended tool with Price Waterhouse for the evaluation of complex customer systems.

"Developing and Deploying Knowledge on a Global Scale" by James Borron, David Morales, and Philip Klahr describes REUTERS, a global-scale help desk that uses case-based reasoning. REUTERS provides a broad variety of new, financial, and other information products in countries around the world, and it is an important component of the business that customers receive fast and effective service when they have difficulty using these services. Reuters undertook to tap the internal expertise of its worldwide organization to build a collection of cases that could effectively provide this support. It built a case base of more than 7700 cases, covering 38 different products and services, that is installed in 29 sites around the world and is used by 190 help-desk personnel who directly service Reuters's customers. The article describes how the company organized and managed the case-acquisition process and tools it developed to help maintain consistency and completeness in the case base.

These five case studies illustrate the broad applicability and the continued maturation of AI technology. AI plays a critical role in the management of business practices, engineering, and even innovative attempts to better connect the U.S. government with its citizenry. These case studies also illustrate another key contribution of AI technology: its ability to help organizations rapidly evolve to meet new needs. This contribution is particularly clear in the White House application: Four years ago, the White House didn't use e-mail, and the World Wide Web barely existed; today, AI technology helps it manage the complexity of the broad-scale use of these technologies.

tion of this research should be related intimately. Our conferences should highlight research that is reaching a level of maturity appropriate for application, and they should also examine the adequacy of our current technology in the context of challenging, innovative applications.

In addition, we would like to create a new focus on emerging technology; *technology* to us means the consolidation of scientific insight in a form that allows routine application. This new focus will include domain analyses, ontologies and core domain theories, paradigm-specific inference engines, and collaborative knowledge-acquisition tools appropriate for the construction of large-scale knowledge bases. In general, this focus will help us examine the insights gained in the attempts to apply AI research and will help us consolidate these insights into useful technology. With the 1996 conference, organizers began the transition to this new view of the IAAI conference. Most notably, we added a series of invited talks that survey areas of AI research that are beginning to emerge into practice, including "AI in Aircraft Design" by Robert Abarbanel from Boeing, "AI in Software Engineering" by Doug Smith of Kestrel Institute, "Speech Technology" by George Doddington of SRI International and the National Security Agency, "Planning" by Mark Boddy of Honeywell, "AI in Computational Biology" by Rick Lathrop of the University of California at Irvine, and "Knowledge Discovery and Data Mining" by Usama Fayyad of Microsoft. These are areas in which we expect to see significant new applications of AI technology in upcoming IAAI conferences. In addition, the IAAI and the AAAI conferences cosponsored an evaluative talk by Frederick Hayes-Roth of Teknowledge entitled "AI, What Works and What Doesn't?"

Organizationally, the IAAI and AAAI conferences are now merged; there is a single proceedings and registration fee, and all sessions are open to all participants, although there are separate calls for papers and different reviewing criteria. The 1997 conference will move further in this direction. The national AI and IAAI conferences will be distinct programs within the common annual meeting of the AAAI. To supplement the award-winning case studies of innovative, deployed applications, IAAI will add a new track of papers on emerging applications, technology, and issues. The acceptance criteria are presented in detail in the call for papers, but the general principle is that application case studies that can teach us some-

thing new about AI technology are the ones that will be presented. Conference organizers are continuing the Innovative Applications Award for applications that meet the criteria established in past IAAI conferences, that is, those that meet the highest standards of innovation, are fully deployed, and exhibit substantial paybacks. We have added a nomination form to the call for papers in the hope of finding a broader range of innovative applications that qualify for the award, especially those that increase representation of the full range of AI applications and research areas. Other applications, which do not meet all the deployment, innovation, and payback criteria but which yield insight about the application of AI technology, will be accepted for the new emerging applications, technology, and issues track of IAAI.

We hope these changes further improve the IAAI conference, broaden the range of people who feel that the conference has something valuable for them, and provide a forum for issues that have found no forum in either the IAAI or AAAI conferences of the last several years. Finally, we ask for the help of all AAAI members in bringing innovative applications to our attention and submitting invited talk nominations and panel proposals to make IAAI serve the role of fostering a mutually productive dialog between fundamental and applied AI.

References

- Rappaport, A., and Smith, R., eds. 1991. *Innovative Applications of Artificial Intelligence 2*. Menlo Park, Calif.: AAAI Press.
- Schorr, H., and Rappaport, A., eds. 1990. *Innovative Applications of Artificial Intelligence*. Menlo Park, Calif.: AAAI Press.
- Smith, R., and Scott, C., eds. 1992. *Innovative Applications of Artificial Intelligence 3*. Menlo Park, Calif.: AAAI Press.

Howard Shrobe is assistant director of Intelligent Systems and Software Technology and chief scientist of the Defense Advanced Research Projects Agency (DARPA) Information Technology Office. He directs programs in the areas of software engineering, information survivability, and AI. Shrobe is on loan to DARPA from the Massachusetts Institute of Technology (MIT), where he has been a principle research scientist in the Artificial Intelligence Laboratory since 1979. From 1982 to 1992, Shrobe split his time between MIT and Symbolics Inc., where he served as a technical director and vice-president of technology. Shrobe received his M.S. (1975) and Ph.D. (1978) from the Artificial Intelligence Laboratory at MIT, where he was a cofounder of the Programmer's Apprentice Project, and his B.S. (1968) from Yale College.