AI Magazine Volume 14 Number 3 (1993) (© AAAI)

The Applied AI Business

The Evolution of Applied AI: A Report on the Fifth Innovative Applications of AI Conference

Sara Hedberg

The Fifth Annual Conference on Innovative Applications of Artificial Intelligence (IAAI) certainly reflected the current state of the commercialization of AI. A new crop of award-winning applications stood as a testimonial to the continuing inroads AI is making into business, saving hundreds of millions of dollars annually. In addition, attendees were introduced to a number of new AI technologies, such as data mining, that are coming out of the research lab and being readied for use.

A new application area, knowledge publishing, came into the spotlight with a Compaq application that has shrink-wrapped its customer-service troubleshooting knowledge for networked printers. It is now shipped with the product and is saving the company an estimated $10 to $20 million annually. Just as help desks were touted as a natural fit for AI technology a few years ago, knowledge publishing is now the next great hope.

Each year, the conference offers several tutorials and panels by some of the most famous and experienced leaders in applying AI to solve business problems, but attendance at the conference continues to decline, from 500 the first year to half that in 1993. The number of papers submitted has also declined. Is this decline the harbinger of doom for AI? I think not. I believe it is more a sign of the (downsizing) times and the need for increased visibility for the conference. I saw many good signs at the conference that applied AI is alive and healthy.

Winning Applications

The first such sign that AI is continuing to be applied to strategic business problems with impressive bottom-line results is the 17 winning applications themselves (each is outlined in this article). These applications were developed by companies such as Compaq, IBM, Ford Motor Company, AT&T, and Boeing. They span the applications' spectrum, from automated software engineering and intelligent front ends for legacy COBOL systems to customer-service and product-configuration applications. The Compaq application (see Customer Service and Sales) was the star of the show.

The manufacturing domain hosted nine of the applications, with one-third of these being for product quality management. Once again, knowledge-based software (KBS) vendor Inference dominated the conference, with 10 of the 17 award-winning applications using its products. Of the other seven, one used the AION DEVELOPMENT SYSTEM, two government applications used CLIPS, one used OPS-83, and three used home-brewed software or did not specify what they used. Application delivery hardware included everything from remote modem dial-in to mainframes to MACINTOSH POWERBOOKS on board the space shuttle.

It is interesting to note that other AI techniques, beyond traditional representation and reasoning, are beginning to appear in production systems. In particular are multiple uses of fuzzy logic, case-based reasoning, knowledge acquisition, and certainty techniques.

A quick glance at the winners this year shows the savings and the tremendous number of workers' jobs being touched by AI technology. Remember, these are only the winners. There are countless applications throughout business and government that are not publicly discussed for reasons of competitive advantage. As such, these applications offer compelling evidence of the growing strategic significance of AI. (The complete conference proceedings is available from the AAAI Press, 445 Burgess Drive, Menlo Park, CA 94025.)

Customer Service and Sales

Compaq and Inference, QUICKSOURCE: This application was the star of the show. Compaq is now shipping its customer-service troubleshooting knowledge for networked printers to customers ordering the product. Using Inference's CBR EXPRESS, it is one of the first examples of "knowledge publishing," a significant future direction for AI. This application is a second-generation IAAI conference winner. Compaq won an award last year for its SMART system, a customer-service representative's assistant. The company has now gone one step further, shrink-wrapping this knowledge of 600 cases and shipping it on diskette with product, saving Compaq an estimated $10 to $20 million annually in customer-service operations alone. It is reducing customers' calls for help by 20 percent. The company believes additional savings can be attributed to marketing and training benefits.

AT&T, Master Production Scheduler Assistant (MPSA) and Customer Order Loading Expert System (COLES): These sibling systems provide intelligent front- and back-end processing of customer order and inventory data stored in legacy COBOL systems on a mainframe. The annual savings because of the MPSA system is $200,000 in personnel costs; other benefits include increased product sales from higher customer satisfaction and better use of inventory. In addition, AT&T reports increases in the quality of work produced and job satisfaction. The system runs on a

personal computer (PC) using ART-IM. The COLES system is an intelligent data-entry error handler for customer order processing. COLES is a nice example of software reuse: Much of the design and code was borrowed from MPSA, so the system was up in only a few months.

Software Engineering

Celite and Andersen Consulting, Knowledge-Based Design Assistant (KBDA): As part of an APPLICATION SOFTWARE FACTORY project, the KBDA expert system was implemented to guide software-module design. It was applied to a 250,000 COBOL development project, and significant results were reported: The time to develop one module was reduced from 20 hours using traditional programming techniques to five hours using KBDA. Training time and personnel requirements also showed a significant decrease. For a system traditionally requiring 50 staff members, development work can be accomplished by about 12 persons with KBDA.

Manufacturing

Boeing, Engineering Standards Distribution System (ESDS): ESDS helps design engineers specify manufacturing processes and select materials parts for commercial jetliners. (A single jet has over one million parts.) This is no small application; it includes 50 knowledge bases, covering most jetliner part types. Deployed in 1989, ESDS is used several thousand times each month. Boeing reports significant but undisclosed savings, including time saved by domain experts and workers, a reduction in redesign, and a reduction in wasted materials and manufacturing operations. The system is accessible to thousands of engineers throughout the company on IBM 3090 mainframes, currently under MVS/TSO but migrating to MVS/IMS-DC. It is a rule-based system and includes special knowledge-acquisition methods and a domain-specific declarative language (Prologlike) tailored to the experts.

Ford Motor Company and Inference, Computer-Aided Parts Estimate (CAPE): CAPE generates, evaluates, and costs automotive parts manufacturing plans for Ford Europe. Ford reported that CAPE reduces estimating response time by 50 percent, improves control over tooling costs, saves as much as 30 percent over less experienced estimators, and is an excellent training tool that enables new estimators to be productive earlier. In addition, the authors speculate that there are potential vehicle cost savings. CAPE combines traditional KBS techniques with fuzzy logic.

IBM, Digitized Expert Pictures (DEPICT): DEPICT is an intelligent manufacturing information repository for IBM's Essex Junction, Vermont, semiconductor manufacturing facility. It serves as a resource for process engineers when investigating manufacturing problems, providing access to past problem-solving learning. It presents information in the form of images and text. Deployed in 1990, it uses rules and certainty factors. Prototyped in IBM's ESE, it is implemented in ART-IM and delivered on MVS.

IBM, Quality and Reliability Expert System (QRES): QRES is a soft-

ware shell that implements manufacturing quality and reliability prediction models from the Operations Research discipline. QRES identifies problems in manufacturing production lines and predicts production quality levels for new components and assemblies during the design phase. Benefits include reduced prediction-cycle time (in one case from three months to one month, meaning the product could be announced two months earlier) and an improved link between design and manufacturing (leading to fewer engineering changes and more robust designs). It runs on a PS/2, 386, or 486 using ART-IM.

IBM, Diagnostic Yield Characterization Expert (DYCE). DYCE provides automated data interpretation and process diagnostics for semiconductor manufacturing product quality and yield loss. Diagnosis time has been reduced from several hours to minutes. IBM reports that the motivation for the system was twofold: (1) to aid employees in interpreting overwhelming amounts of test data and (2) to reduce the need for additional staffing. The company estimates hundreds of hours each month of engineers' and technicians' time are being saved with DYCE. In addition, early detection has saved additional products from being manufactured with the defective process—at a cost of several million dollars for each missed problem. The shell includes pattern matching, fuzzy logic, and certainty techniques and uses ART-IM.

CRIM, Pitch Expert for Kraft Mills: CRIM helps mill-site engineers diagnose pitch problems in the paper-manufacturing process. Now in use at 36 mills in Canada, representing 20 companies, current annual savings estimates are $28 million Canadian. The system was designed over 4 years, using ART on a SPARCSTATION 1 and requiring 21 person-years and $2.8 million Canadian to develop. It is now accessible by modem, operating in a question-and-answer mode.

Configuration

AT&T Bell Labs, Product Offerings Expertise (PROSE): Call this product AT&T's E-XCON. PROSE is a general-purpose configuration shell designed to support the sales and manufacturing order process. AT&T expects the following general benefits as PROSE applications spread throughput its operations: a reduction in order-entry errors; staff consolidation; an improved product-design change cycle; and a faster, more direct order-entry process. The three current installations of PROSE run on a SUN 490.

NYNEX, Outside Plant Engineering and Resource Administration (OPERA): OPERA assists outside plant telecommunications engineers with planning and installing sophisticated electronic equipment. NYNEX reports a productivity gain of 3 to 6 times and a $5 million annual savings. It cost under $1 million to develop and was completed in 3.5 person-years. The system interface is a set of schematic documents and forms. It is written in C and OPS-83, integrates with a relational database, runs on a SPARCSTATION 2, and is modem accessible. Deployed in July 1992, it is in operation in 34 engineering centers in New York and New England, assisting over 350 outside plant engineers.

Troubleshooting

General Electric and U.S. Department of Energy, DODGER—A Diagnostic Expert System for the Evaluation of Nondestructive Test Data: DODGER diagnoses heat-exchanger tubing based on enormous quantities of data from eddy current signals. Where only 5 percent of all human checks were double-checked manually in the past, now 100 percent are automatically checked with a high confidence level. DODGER uses more traditional KBS techniques plus machine pattern recognition and fuzzy logic.

Southwest Research Institute and U.S. Air Force Materiel Command, Deficiency Report Analysis Information Reports (DRAIR): DRAIR troubleshooting adviser provides United States Air Force engineers with aircraft parts performance history and recommends improvements. With this system, the Air Force estimates a time savings of 1,900 person-hours each year and increased productivity from reduced staff equaling $120,000. Accessed by PC-modem or Internet, it was built using CLIPS (an ART-like, OPS-like system).

Northrop and Inference, Guidance and Control Expert System Shell (GCESS): GCESS is an application shell for symptom-driven diagnosis that has been used with Peacekeeper missiles for maintenance fault isolation. It was developed with ART-IM and runs in Microsoft WINDOWS. The cost justification was based on operational necessity, not cost-benefit.

Law Enforcement

Tennessee Department of Correction and Andersen Consulting, Tennessee Offender Management Information System (TOMIS): TOMIS integrates KBS to calculate offenders' sentences within a large COBOL-DB2–based offender management system. The KBS module has increased sentencing accuracy from 80 to 100 percent and saves more than 37,500 person-hours of manual calculation each month. It also enables less experienced workers to calculate sentences. A workstation version will be used to test changes to laws and guidelines, allowing the state to analyze the impact of changes prior to enactment. Developed in Trinsic's AION DEVELOPMENT SYSTEM, it runs overnight in batch mode on a mainframe.

Space Research

NASA-Ames, Principal Investigator (PI)-IN-A-BOX: AI is no longer pie-in-the-sky whimsy judging from this PIE-IN-A-BOX from the National Aeronautics and Space Administration (NASA). This application assists astronauts in performing a variety of scientific experiments in space. It includes several modules to assist with acquiring and evaluating data, diagnosing and troubleshooting data-collection problems, and identifying and scheduling the best experimental protocol. NASA estimates a 20-percent increase in crew productivity for each astronaut science-hour with PIE-IN-A-BOX. It runs on a single MACINTOSH POWERBOOK 170 (8-megabyte [MB] random-access memory, 40-MB internal hard drive), with an external analog-to-digital

converter connected to the POWERBOOK's SCSI port. It was implemented using CLIPS and HYPERCARD.

Informative Tutorials

The IAAI tutorials on Sunday included an impressive array of the world's most experienced AI practitioners. Attendees had access to AI applications leaders such as Ed Mahler and Lynden Tennison, who told real stories, with humor and elan, of their experience in the business world.

The "AI in Customer Service" tutorial, for example, really benefited from the humorous anecdotal teaching style of Avron Barr (Inference), one of the first to recognize and realize the potential of the help desk–AI marriage. Barr's theoretical AI underpinnings and his years of application experience provided a rich source of information for attendees interested in actually building help-desk applications. Anil Rewari of Digital Equipment Corporation, who co-led the customer-service tutorial with Barr, outlined several promising research areas for help-desk applications: knowledge sharing, natural language processing, and machine learning. Thus, attendees were able to get a good picture not only of how today's technology can help automate the help-desk functions but also of what powerful AI paradigms are coming down the pike.

Another sign of the growing maturity of applied AI was to be found in the business process reengineering (BPR) tutorial. In the early days of applied AI, there was little sensitivity to business problems. Today, it is clear to all that this area must be the starting point for any AI application. In the BPR tutorial, for example, Neal Goldsmith, champion of the famed American Express AUTHORIZER'S ASSISTANT and now president of Tribeca Research, gave a good crash course in the theory and practice of BPR followed by some Inference examples of AI BPR projects.

Panels Look to Emerging Technologies

The "Emerging AI Technologies Panel," organized by Neena Buck, had a blue ribbon group of AI practitioners who talked about new ways that AI will make computers smarter. Buck, formerly with New Science Associates and now with Flat Earth Systems Group, organized a thought-provoking panel that showed the audience what's emerging from the research labs. Emerging technologies that were touched on include massive parallelism, fuzzy logic, genetic algorithms, case-based reasoning, natural language, speech recognition, and neural networks.

Both Liz Byrnes (Banker's Trust) and Ron Day (Arco) expressed the need for better software tools with small modules that are callable as needed. Day stated that users need to be able to mix and match technologies without the extra baggage that comes with each vendor's tool. A new offering from Neuron Data, called SMART ELEMENTS, which was demonstrated at the National Conference on Artificial Intelligence vendor show, was an example of a tool that could meet this need.

From almost every panelist, we heard mention of "database mining" techniques, which use natural language methods to find those nuggets embedded in the vast caverns of legacy databases. This emerging technology, which most panelists felt is still in the research stage, holds tremendous promise for businesses. Another emerging area, identified by several as critical for future knowledge distribution, was the sharing of knowledge between different knowledge bases written with different tools.

According to the "Next-Generation Knowledge-Based Systems" panel, organized by Herb Schor of USC/Information Sciences Institute, help is on the way for sharing knowledge. Rich Fikes (Stanford University), a leading knowledge representation scientist, described the knowledge interchange format (KIF) work under development courtesy of the Defense Advanced Research Projects Agency (DARPA). This seminal work could have a profound impact on the problems of addressing heterogeneous applications and languages—a growing problem as corporations become

increasingly linked by computer networks. Fikes and his associates are developing an interlingua to serve as an intermediary between systems.

I asked Bill Marks of Lockheed, whose group is a test bed for this technology, when the technology would be mature enough to cross over from the research labs into general use. Marks replied that a good deal of the functions were already available and that an important next step is wider use and testing.

Beyond the Human Metaphor

There was much talk throughout the conference regarding what insiders perceive as the AI identity crisis. For invited speaker Joe Carter of Andersen Consulting, the solution starts with our perceptions of AI. Carter, who heads one of the premier AI efforts in the world, asserted that "AI set out to build a telegraph and ended up with a radio." This telegraph, according to Carter, was the human metaphor—the emulation of humans in computer systems. The radio that we ended up with is, in Carter's experience, the dozens of applications his consulting group has built that do not emulate the human decision-making process at all. In fact, they improve the process.

By way of example, he cited an airline auditing system his group built to ensure that fares were applied properly, appropriate taxes collected, and agent commissions taken appropriately. "Our system changed the way the airline performed the auditing. They used to take statistical samples. Now they process all tickets."

By looking only for experts, Carter believes we have made a mistake. We must move beyond the classic notion of the expert system. In fact, the expert can be more of a constraint than a model. To underscore this point, Carter told about an application his group built for an Australian insurance company that encoded an expert's underwriting heuristics. Upon examination, the group realized that he applied rigorous standards to men driving VWs. After digging a bit, it was found that the expert's wife had left him for a man who drove a VW!

Carter gave the audience a call to action. "It's time to take a Copernican view of intelligence and knowledge. Humans are not the sole source of knowledge or intelligence. If we think so, then we're doing a disservice to AI technology. ... The human metaphor is too limiting. The public view of AI—that we're trying to emulate humans—is rendering misguided expectations of AI. We must get people to unlearn the human metaphor."

Evolution Takes Time

Overall, I believe that we all need some perspective and patience when we evaluate the progress of applied AI. Advanced technologies seldom take the world by storm overnight. When the Wright Brothers finally got humanity airborne, it was in a glider with a 12-horsepower engine, not in the space shuttle. The shuttle took the better part of the twentieth century.

… by all indications, work is continuing, and applications are multiplying…

It is clear that people outside the industry have misconceptions about AI. With IAAI continuing to be a focal point for the evolution of applied AI, with time this misconception will be replaced with a truer view of the value AI has for business today and in the twenty-first century. Let's not disparage the current state of applied AI. AI will not disappear. Much like the Roaring 20s, we had the 1980s: the lavish parties, the initial flurry of money, the media attention. We're now in a more sober economy of diminished resources, yet by all indications, work is continuing, and applications are multiplying.

I believe Ed Feigenbaum's talk underscores this point. He examined the question of why AI is like a tiger in a cage. He pondered why other technologies, perhaps less powerful (such as relational databases) have experienced exponential growth when AI seems stalled. I won't give a full description of his talk, but Feigenbaum made a compelling case for the continued contribution that AI applications are making. Perhaps most telling was the DART application that scheduled the movement of all personnel and materiel to Saudi Arabia for Desert Storm. This application, according to the director of DARPA, Victor Reis, in a talk at Stanford, paid back DARPA's entire investment in AI over the years. (Reis did not specify whether this return was in dollars or some other form.) Here is yet another sign of the continued significance of the technology.

Feigenbaum made another important observation during his speech: "AI technology is ahead of the market. Desktop power with large memory, GUIs, object-oriented programming networks, and rapid prototyping (long-time standards of AI programmers) are just now part of MIS." Change aided by complex tools takes time; so, let's give applied AI time to evolve—because it will. The business benefits are too compelling for it to do otherwise.

Next Year

I'm looking forward to next year's conference in Seattle, Washington. IAAI continues to be the place to see the latest application breakthroughs and find out what's coming out of the AI research pipeline.

Sara Hedberg is president of Emergent Inc. (Issaquah, WA), a company specializing in analyzing emerging trends in software technologies and products. She has been an AI industry analyst for six years, has written recent articles for *Byte* magazine on AI trends, and has served as the editor-in-chief of the *Spang Robinson Report on Intelligent Systems*. Hedberg has worked in the AI industry since the early days of commercialization in 1983. Hedberg can be reached at hedberg@halcyon.com.