

Commercial AI Trends Seen at AAAI-87

Jeffrey Stone

The annual conference of the American Association for Artificial Intelligence (AAAI) is the largest and most important meeting of AI theoreticians and practitioners in the United States. This year, the conference was held in Seattle, Wash., and paid attendance was just under 5100. Last year's Philadelphia conference drew 5400. The drop in attendance was primarily the result of competition with the International Joint Conference on Artificial Intelligence, which took place in Milan a few weeks after AAAI.

Although the AAAI conference technical program caters to the academic and theoretical side of AI, the conference exhibits and related vendor presentations provide a unique insight into the commercial status of AI. I took advantage of the AAAI-87 exhibits to visit with a number of industry leaders and discuss the state of the U.S. AI marketplace.

The State of the AI Marketplace

Vendors everywhere tend to be optimistic about their products and industries. Certainly, the AAAI vendors displayed optimism, but there was also an underlying current of doubt—or at least question. Questions raised by the vendors include the following: Does the recent lull in the market mean that the luster of AI and expert systems in particular is wearing thin? Are there more than a handful of significant expert systems in operation today? Are symbolic machines and high-end complex expert system shells on the way out? Is expert system technology just another software methodology that will soon blend into the software landscape?

No simple answers exist to these questions, but I discerned a number of themes that help put the marketplace in perspective.

Theme #1:

The Evolution of the AI Marketplace

Until the late 1970s, AI was limited primarily to research activities in university laboratories. In 1980, the first Lisp machines appeared in the commercial marketplace. Most of the action remained in hardware through 1984. The first full year in which

expert system software played a significant commercial role was 1985.

The introduction of tools such as KEE from Intellicorp in 1983 and S.I. from Teknowledge in 1984 opened the marketplace by drastically reducing the amount of effort and development time necessary to build expert systems. Beginning in 1984, a group of *early adopters*, companies willing to lead in the exploration of expert system technology, procured hardware and software tools.

One of the messages from the early adoption phase was that expert systems were fine vehicles for fast prototyping and, ultimately, for rapid application development. This message was strongly broadcast throughout the information system industry and created some unrealistic expectations.

A slowdown in the growth of the demand for expert system hardware and software has recently occurred. This softness is partly the result of the same trends that have slowed the entire information systems industry since late 1985. In addition, however, demand for expert system products has lagged because the early adopters have not yet deployed applications. This situation has two important consequences: (1) Early adopter organizations still in the development phase have not yet procured all the **operational** tools that will eventually be required. (2) Organizations that were not early adopters (the larger group) are waiting for success stories before taking the plunge. These success stories will come when the early adopters deploy major applications.

Theme #2:

Few Major Applications Currently Deployed, Many More on the Way

Most industry insiders believe that

until now, only a handful of expert systems with strategic organizational impact have been deployed. Examples of this sort of system are American Express's credit limit authorization application used by several hundred operators and Digital Equipment Corporation's suite of expert systems that support VAX sales and manufacturing.

Although few such applications are now in operation, many more are in

growing demand for expert system consulting services. As organizations move closer to deployment of their first operational application, there is increased sensitivity to the fact that the organization is betting a great deal on new technology which is yet unproven. Managers are highly motivated to seek outside aid to solve remaining problems and to help provide assurance that the applications

stories of the early adopters.

The second attribute is rapid payback. Managers want to justify their decisions, and the most frequent justification for the capital expenditures associated with an AI project is rapid payback in the form of operational system savings. According to an IBM AI marketing manager, among IBM's customers, the organizations that are currently deploying are those where the chief financial officer can foresee a solid return on the investment one or two years after deployment.

Where AI technology was once acquired by research and development laboratories with "easy money," procurement now requires a traditional sales cycle that starts at need identification and stresses cost justification.

... the organizations that are currently deploying are those where the chief financial officer can foresee a solid return on the investment one or two years after deployment.

the final stages of development. Within the next 12 months, the number of deployed strategic-value expert systems is expected to jump from a few dozen to well over a hundred. This expansion will provide the second tier of organizations (the non-early adopters) with enough ammunition to begin serious implementation of their own applications.

Theme #3:

Applications Becoming Less Complex

When AI first moved into corporate laboratories, the projects were large and ambitious. Many research expert systems grew knowledge bases with more than 1000 rules; systems with over 500 rules were the norm. However, when it came to thinking about operational systems, managers adopted a much more conservative position. The systems that are already operational or will soon become operational are smaller in size and less ambitious than their forerunners. Rather than the ability to support complex problems, the emphases for operational applications are low risk, high payoff, and smooth integration with conventional computing environments.

Theme #4:

Demand for Services Growing

Another effect of the shift from laboratory to operational application is a

will succeed.

Most vendors I talked with have already created or are now in the process of creating an organization that provides customers with consulting services. Several vendors said the demand was currently so large that they would be able to service only a portion of the potential business.

Theme #5:

AI Is Now Acquired through the Traditional Procurement Cycle

Just a few years ago, AI was viewed by most managers as an esoteric technology that had little place outside the laboratory. Things have changed significantly, however, and today, there is widespread recognition on the part of U.S. management of the potential benefits of AI and, specifically, expert system technology. Although such recognition was enough to bring AI into corporate laboratories, it is not sufficient to bring it into corporate operations. Managers today are looking for two additional key attributes.

The first of these attributes is prior success, which was already mentioned. The way to understand this need is to understand that the second tier of organizations, those which have not yet adopted the technology, will not do so until they feel secure that their programs have a reasonable chance for success. The only real source of such security is the success

Theme #6: Integration with Conventional Data Processing Environments

Integration with conventional data processing environments was the dominant story on the AAAI-87 exhibit floor. It seemed as if all the vendors had been given a training class where they learned how to talk about delivery of expert system applications in mainstream corporate computing environments. Some of the vendors even took the advanced training class where they learned about developing systems in these same environments.

Regardless of how the vendors learned their lessons, they all seemed to tell essentially the same story when it came to integrating AI into the corporate computing environment: (1) Applications are delivered running on the same systems that are used to provide mainstream corporate applications, which usually means mainframe computers and minicomputers from vendors such as IBM and Digital. (2) Applications need full access to standard database systems running in their target environments. Full access involves the ability to write, as well as read, database records and manipulate these records as consolidated objects rather than as loose collections of data items. (3) Users want to be able to embed an expert system subsystem within existing applications. The majority of significant applications now being developed

involve expert system enhancement of existing strategic applications. The ability to embed is critical.

Theme #7: User Training Is Essential

One of the constraints that is limiting the penetration of expert systems is the small quantity of capable knowledge engineers in user organizations. Without solid training capability, vendors working with large organizations find limited acceptance of their expert system products. For vendors dealing with such organizations (for example, Digital, IBM, Inference, Intellicorp, and Teknowledge), user training is a high priority. There is no shortcut or easy way through it. Generally speaking, users need training and support until they can complete a meaningful pilot project. Once there are individuals who have completed an initial pilot project, other people within the organization can be brought up to speed with the technology.

Theme #8: Shell Pricing

The current price point for high-end expert system shells is somewhere in the \$40K to \$50K price range. Intellicorp's KEE is the leading example. Discussions at AAAI-87 led me to believe that KEE and other high-priced products will not be able to maintain their premium pricing in the face of challenges from products such as Gold Hill's GoldWorks and Neuron Data's Nexpert. These products provide as much as 90 percent of the functionality of the premium-priced products but are priced well under \$10K.

Products such as KEE, however, will always be attractive for users with large applications whose development costs run into the millions of dollars, but such applications are relatively few. For the majority of applications, 90 percent of the functionality at 20 percent of the price is most attractive. Ultimately, I expect to see companies with high-end expert system shells, such as Intellicorp, develop new products that are priced much lower and can compete with the Nexperts and the GoldWorks. The premium products will be maintained for the relatively few applications where tool price is not a major consideration.

Market Research Firm	HW 1987	SW 1987	Total 1987	HW 1990	SW 1990	Total 1990
IDM Data	710	845	1555	1700	2500	4200
Frost and Sullivan	274	665	939	740 ¹	527 ¹	1267 ¹
IDC	341	297	638	440	866	1306
Input	NA ²	NA ²	350	NA ²	NA ²	885
Macintosh International	245	250	495	511	455	966
Ovum	227	197	424	643	519	1162

¹ This estimate is for 1989 rather than 1990
² NA - Not available

Source: Numbers are taken from Walker, F. C. and Miller, R. K. April 1987 *Expert Systems 1987: An Assessment of Technology and Applications*. Madison, Georgia: SEAI Technical Publications

Table 1. AI Market Estimates and Projections (market values in \$ millions)

Marketplace Status Summary

AI, specifically expert system technology, has gone through an important transition. The primary corporate interest in the technology is no longer research oriented. Interest stems from the desire to deploy applications that utilize the technology in operational settings. The number of strategic applications, those making significant contributions to an organization's bottom line, is beginning to grow and is expected to grow even more rapidly over the next few years.

The transition from research to operational focus brings with it a completely new set of user requirements. Where in the past the greatest search scheme and the most sophisticated development interface were hot items, today the emphasis falls to more practical considerations: integration with, and delivery into, conventional computing environments; cost justification, usually by way of rapid payback; enhanced levels of vendor support both in training and consulting; and less complex applications with greater likelihood of success.

The bottom line is that the commercial market for AI is no longer technology driven as it has been in the past. It is now driven by the very real

needs of users who deploy the applications upon which their organizations depend. The new realities of life indicate that vendors which have stressed the practical business aspects of AI deployment will embark upon the next phase of marketplace development with considerable advantage.

Market Estimates

My view of the commercial AI market is that to date there has been a lot of "tire kicking" and that we are just now seeing the first crop of significant operational applications. Over the next few years, I expect to see considerably increased activity as large corporations deploy strategic systems that depend on expert system technology. This projection is supported by market estimates from most of the research firms following the AI marketplace (see Table 1).

Final Word

The most concise summary that I heard of the state of commercial AI as seen through the AAAI exhibits is a little gem from Harry Reinstein of Aion: "AI is no longer a technology-driven market, it's a market-driven technology."

Errata to the Fall Issue (Volume 8 Number 3)

Page 42: For d_1^2 read d^2

Page 44: For β might be taken ... read β might be taken

Page 49: For $\$2_{n-1}$ read $\$2^{n-1}$