AI Magazine Volume 15 Number 2 (1994) (© AAAI)

Photo courtesy NKK Corporation

NKK Steel Blast Furnace, Kawasaki, Japan.

An expert system used in the control room of this blast furnace controls fluctuations in furnace temperature, thereby saving significant amounts of energy and costs. (See story.)

Knowledge-Based Systems Research and Applications in Japan, 1992

Edward A. Feigenbaum, Peter E. Friedland, Bruce B. Johnson, H. Penny Nii, Herbert Schorr, Howard Shrobe, and Robert S. Engelmore

■ This article summarizes the findings of a 1992 study of knowledge-based systems research and applications in Japan. Representatives of universities and businesses were chosen by the Japan Technology Evaluation Center to investigate the state of the technology in Japan relative to the United States. The panel's report focused on applications, tools, and research and development in universities and industry and on major national projects.

In 1991, the National Science Foundation and the Defense Advanced Research Projects Agency tasked the Japan Technology Evaluation Center (JTEC) with studying the state of knowledge-based systems research and applications in Japan. JTEC formed a panel of individuals from the academic and business communities to conduct this study.

The primary objectives of the JTEC panel were to investigate Japanese knowledge-based systems development from both technological and business perspectives and compare progress and trends with similar developments in the United States. The panel focused on (1) applications in the business sector, (2) infrastructure and tools, (3) advanced knowledge systems development in industry, (4) advanced knowledge systems research in universities, and (5) national projects.

The JTEC panel visited 19 sites during its 1-week visit to Japan in March 1992 and conferred with Japanese computer scientists and business executives both before and after the official visits. The panel visited four major computer manufacturers; eight companies that are applying expert systems to their operations; three universities; three national projects; and *Nikkei AI,* a publication that conducts an annual survey of expert system applications in Japan. This article summarizes the findings of the panel in each of the five areas listed.[1]

Members of the Panel

The panel members were Edward Feigenbaum, Stanford University (chair); Peter Friedland, National Aeronautics and Space Administration; Bruce B. Johnson, Andersen Consulting; H. Penny Nii, Stanford; Herbert Schorr, University of Southern California; and Howard Shrobe, Massachusetts Institute of Technology and Symbolics, Inc.). Robert Engelmore served as an ex officio member of the panel with the responsibility of producing the final report. Also present on the site visits were Y. T. Chien, National Science Foundation, and R. D. Shelton, JTEC.

Sites Visited

The sponsors of the JTEC study defined the dimensions of the study to include the following areas: (1) business-sector applications of expert systems; (2) advanced knowledge-based systems in industry; (3) advanced knowledge-based systems research in universities; (4) work at government laboratories, especially the laboratory of the Japanese Fifth-Generation Computer Project; and (5) the electronic dictionary research knowledge base building effort. The panel was also asked to observe the fuzzy system work being done in Japan, any neural network applications that affect expert system development, and the new national project known as Real-World Computing.

Half of this study effort was aimed at applications of expert systems in the business sector. Knowledge-based systems research in industry made up 15 percent of this study and knowledge-based systems research in universities another 15 percent. Two national laboratories (Fifth-Generation Project and Electronic Dictionary) each account for 5 percent of the study. The remaining 10 percent focused on miscellaneous topics.

During the week of our visit, 19 visits were made. Applications of expert systems to business-sector problems and the industrial knowledge-based systems research together accounted for 12 of the 19 visits. University knowledge-based systems research accounted for three visits, ICOT and EDR accounted for two, and other visits accounted for two.

We chose the industrial sites to be visited based on the following criteria: (1) computer-manufacturing companies that were known to be active in knowledge-based systems applications and research (sites visited: Fujitsu, Hitachi, Toshiba, NEC); (2) noncomputer companies at which there was at least one well-known expert system application (sites visited: Japan Air Lines, Sekisui Chemical); and (3) selected companies from certain industry groups that were known to be active and highly competent in building expert system applications, for example, steel, construction, electric power, and communications industries (sites visited: NTT, Tokyo Electric Power, Nippon Steel, NKK Steel, Mitsubishi Electric, Obayashi).

The panel saw strong efforts by Japanese computer companies and industry-specific companies (for example, Nippon Steel) to advance their knowledge-based systems technology and business.

Visits with university professors were decided on based on personal knowledge of the leaders in academic knowledge-based systems research in Japan. As it happens, these leadership positions were held by professors at the major universities: Tokyo University, Kyoto University, and Osaka University.

Finally, we scheduled a special visit with the editor and the staff of *Nikkei AI* newsletter to check facts that we had accumulated and impressions that we had. *Nikkei AI* is the leading Japanese news publication in the field of knowledge-based systems applications and research.

General Observations and Conclusions

There are many similarities between Japan and the United States in both research and application activities. However, there are some noteworthy contrasts, which are highlighted in the following paragraphs.

Japanese computer manufacturers play a dominant role in the technology and business of expert systems. These manufacturers have mastered and absorbed expert system technology as a core competence. They tend to use system engineers, rather than knowledge engineers, to build systems. Consequently, integration with conventional information technology poses no special problem and is handled routinely and smoothly. These large computer companies also build many application systems for their customers; smaller firms specializing in AI software play only a minor role in application building, in contrast to the situation in the United States.

Japan has more experience than the United States in applications of knowledge-based systems technology to heavy industry, particularly the steel and construction industries.

Products based on the use of fuzzy control logic have had a big impact on consumer products, including video cameras, automobile transmissions and cruise-control systems, televisions, air conditioners, and washers and dryers.

The panel saw strong efforts by Japanese computer companies and industry-specific companies (for example, Nippon Steel) to advance their knowledge-based systems technology and business. This situation contrasts with that in the United States, where there is a declining investment in knowledge-based systems technology: lack of venture capital, downsizing of computer company efforts, and few new product announcements. It is a familiar story and one of concern because this trend can lead to Japanese superiority in this area relatively soon.

Although the quality of research at a few top-level universities in Japan is in the same range at top-level U.S. universities and research institutes, the quantity of research (in terms of number of projects or number of publications) is considerably smaller (by nearly an order of magnitude).

Specific Findings

Specific conclusions and observations for each of the study areas are reported in the following subsections.

Applications of Knowledge-Based Systems in Japan

Each year, *Nikkei AI* publishes a survey of expert systems in Japan. For its 1992 special issue, questionnaires were sent to 2200 companies, of which 295 responded. Figure 1 shows a steady, approximately linear growth in the number of systems, both fielded and under development since 1987. The rate of increase in the number of companies that are

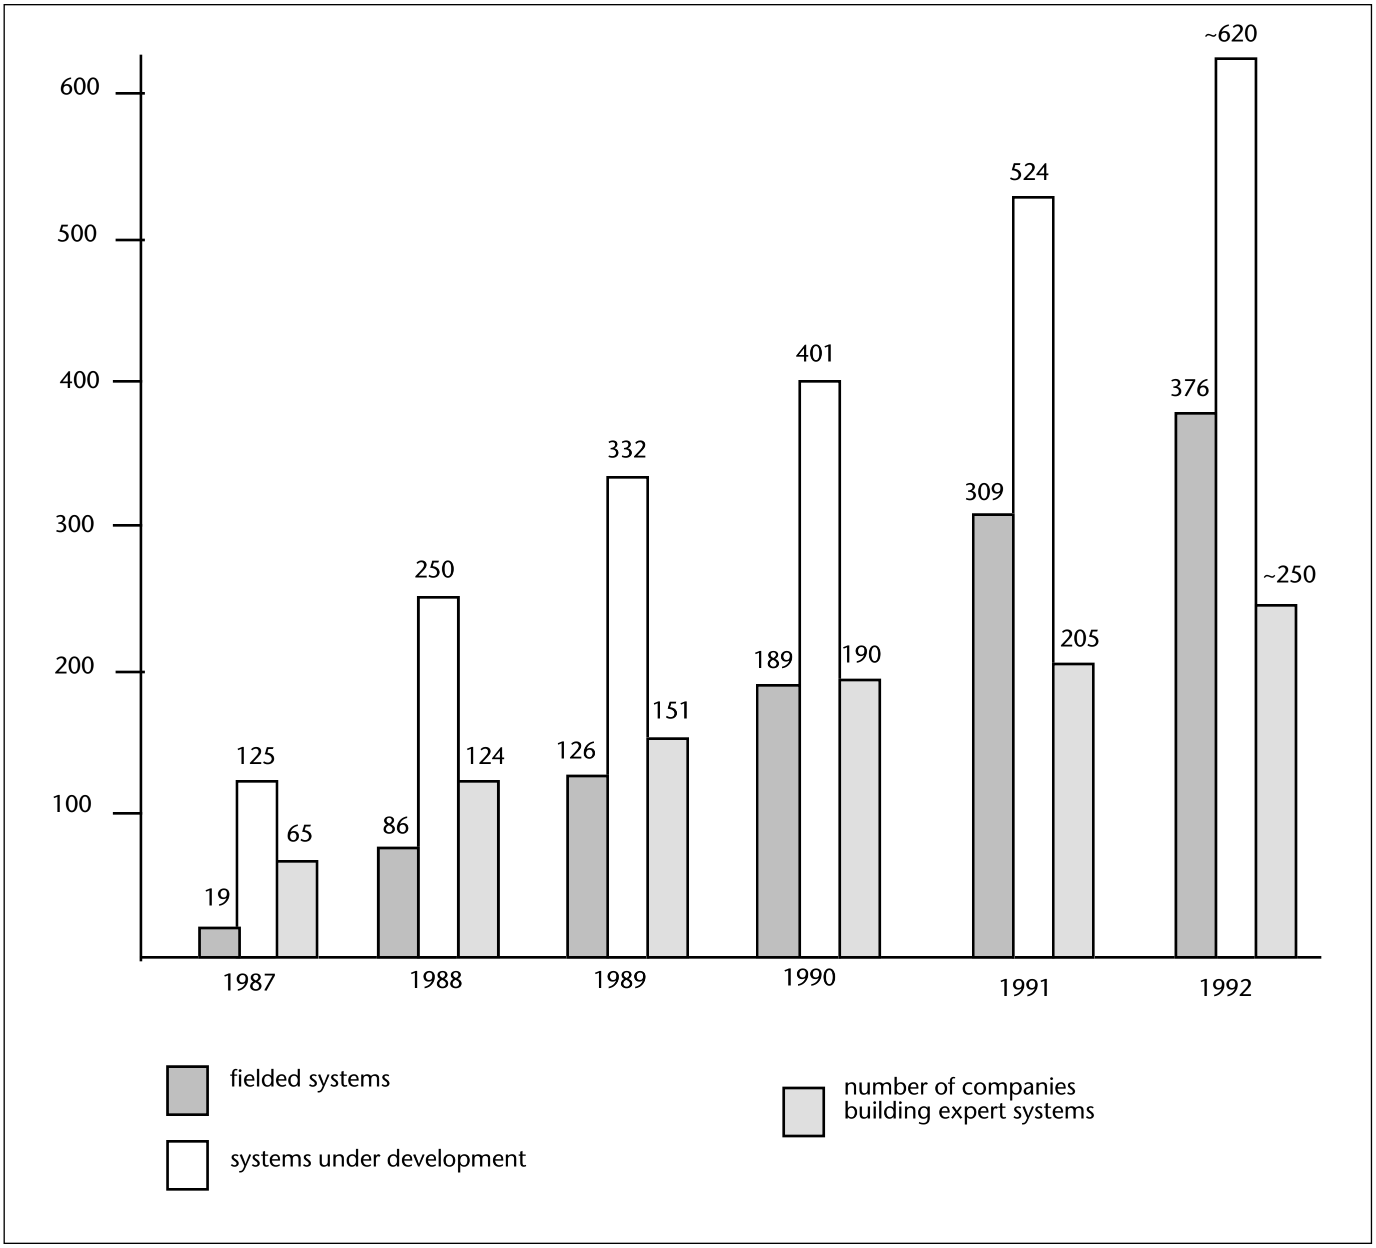

Figure 1. Growth of Expert Systems in Japan (Nikkei AI 1992; used with permission).

developing or using the technology has slowed in recent years. Although we are aware of no comparable study for companies in the United States, we would expect a similar growth pattern.

The most prevalent types of application that have been fielded are for diagnosis, planning, and design. Diagnostic systems, which were the most numerous (over 40 percent of the total), include help-desk services and management analysis aids.

The *Nikkei AI* survey also looked at expert system development from the viewpoint of organizational structure, that is, who actually builds the systems. Here, we see some contrast with the situation in the United States. In Japan, most expert systems are built by an organization's management information system group or by its research and development division. Relatively few expert systems are contracted out for building. In the United States, we suspect that a much larger percentage of expert systems are built by the numerous outside companies that specialize in building custom systems (again, no domestic survey has actually been done). The larger

These knowledge systems reduce the cost of making modular houses, improve the quality of their products and services, and reduce the error rate. The steady-state 5-percent error in part selection has been reduced to near 0 percent.

Japanese companies appear to have a tighter coupling among their research and development, product development, and operations divisions, so that internal development and deployment of new technology is more easily achieved. Moreover, according to *Nikkei AI,* the trend over recent years has been an increased migration of application development from the laboratories to the information system divisions, indicating that the expert system technology is now in the mainstream.

Some other general conclusions from the recent *Nikkei AI* survey that we found support for during our visits include (1) a growing interest and involvement in related technologies, for example, fuzzy logic, neural networks, and object-oriented databases; (2) a steady move away from specialized computers (for example, Lisp machines) toward UNIX-based workstations; (3) a recognition of the value of knowledge-based systems in improving individual decision making, working procedures, and timeliness; and (4) a recognition by some companies of operational problems, such as obsolescence, maintenance, and completeness of the knowledge base. (One would expect similar problems with database systems.)

Case Studies: High-Impact Systems

The importance of a new technology such as knowledge-based systems can be measured by the depth and the breadth of its impact. In the following subsections, we discuss the breadth of the knowledge-based systems applications that we observed in Japan. We provide some insight into the impact of knowledge base systems with three examples of systems that appear to have made significant changes in the way in which organizations conduct their daily business. In all cases, knowledge base technology was used to augment and replace traditional computational technology for reasons of speed, reliability, and quality of solutions.

Modular House Configuration (Sekisui Heim) In 1971, Sekisui Chemical created the Heim division to build modular houses. Unlike prefabricated houses, modular houses are semicustom houses designed and constructed out of modules. More than 80 percent of the house is built in one of Sekisui's six highly automated factories. Sekisui Heim, currently the fourth largest house builder in Japan, builds about 20,000 houses each year.

After a house is designed, all the necessary parts must be identified and delivered to the factory floor in a timely manner. An average house requires about 5,000 unique parts; Sekisui has 300,000 different parts in stock. Prior to the installation of an expert system, every time a new house design was introduced, there was an error rate in part selection of about 30 percent for the first 6 months. As time went on, the error rate decreased to about 5 percent. The high initial error rate made the introduction of new products highly problematic, and the steady-state error rate of 5 percent cut into profits.

Sekisui acquired a majority interest in a startup AI company, which is developing a number of knowledge-based systems for the parent company. The first expert system to be put into routine use, HAPPS, identifies and selects the necessary parts for a given house design and schedules the delivery of the parts to the right place on the factory floor at the right time. Two other systems for other styles of modular houses have since been developed and put into routine use. These knowledge systems reduce the cost of making modular houses, improve the quality of their products and services, and reduce the error rate. The steady-state 5-percent error in part selection has been reduced to near 0 percent.

The three systems cost approximately ¥450 million (US$3.5 million) to build. The development cost is expected to be paid back within two to three years. Sekisui claims that the return on investment has been ¥1 billion (US$8 million) annually.

Aircraft Crew Scheduling (Japan Airlines) Our next example comes from Japan Airlines, Japan's official international airline (although it also does significant domestic business), which developed a knowledge base system for crew scheduling (Onodera and Mori 1991). The airline maintains a fleet of over 100 wide-body aircraft: Boeing 747, 747-400, and 767 and McDonnell-Douglas DC-10. It has a staff of about 2200 flight crew members (pilots, copilots, and flight engineers). The problem is to produce a monthly crew-allocation schedule that takes into account a wide variety of constraints. These constraints include crew training (for aircraft and route qualification), restriction on maximum number of takeoffs and landings, vacation and meeting needs, and crew turnaround times at various destinations. The schedule for a given month needs to be produced by the twenty-fifth of the preceding month to give adequate warning to crew and maintenance personnel.

Before the development of the knowledge base scheduling system, called COSMOS/AI, about 25 human schedulers were involved in solving the problem. The hardest schedule (for 747s) took 20 days (with a great deal of overtime) to prepare. Moreover, the schedulers needed about a year to become expert in the problem. An important, related issue was the maintenance of scheduling knowledge, that is, updating information on planes, crews, government regulations, and so on. In the summer of 1986, the airline decided to investigate various automated approaches to improving their solution to the crew-scheduling problem. It developed two automated approaches to the problem, one a traditional operation-research scheduling system (in cooperation with Hitachi) and the other a knowledge-based approach (with NEC).

Testing of both systems began in the summer of 1988. The knowledge base system was easily the system of choice for two major reasons: First, it produced better, less mechanical schedules because it was far more able to represent complex yet informal constraints on crew preferences and other related factors; second, it was much easier to maintain. The knowledge base system became fully operational in February 1990. It has reduced the number of human schedulers from 25 to 19 while airline operations have increased by 5 to 10 percent. The aforementioned 747 schedule now takes a maximum of 15 days to produce, with no overtime required, compared to 20 or more days, including overtime, to do the job previously. Training time has been reduced to two months. Overall, scheduling productivity has approximately doubled. The system cost about ¥500 million (US$3.5 million) to build, and it paid for itself in direct cost savings in about 18 months. Japan Airlines views the harder-to-measure savings in increased crew satisfaction and ease of maintenance as equally important.

Blast-Furnace Control (NKK Steel) The production of pig iron in a blast furnace is a complex, distributed, nonlinear process. Conventional mathematical modeling techniques have never been able to predict future states of the furnace with enough accuracy to support automated control. As early as 1986, NKK Steel's Keihin Works had developed an expert system to predict abnormal conditions within the furnace. The system became operational in 1987. NKK and other Japanese steel companies have since developed other knowledge-based blast-furnace–control systems.

The challenge for the furnace controller is to minimize the uncertainty in the operating temperature. The smaller the uncertainty, the lower the overall temperature is that is needed to produce the pig iron (see figure 2), resulting in large fuel savings.

An expert system has been developed that successfully models the current state, predicts future trends with sufficient accuracy to make control decisions, and actually makes the control decisions. These decisions can be implemented automatically, or the operator can take manual control but still operate through the expert system's interfaces.

The blast-furnace–control application is noteworthy for many reasons. The problem had not been solved previously by other techniques, and the system was developed by system engineers, not knowledge engineers. It is in daily operation now at two plants and will soon be installed in two more. The company reports an estimated annual savings of approximately ¥800 million (US$6 million), a reduction in staff of 4 people, and an improvement in the quality of the furnace output because of reduced fluctuations in furnace temperature.

The expert system itself is an integral part of the total furnace-control system, which makes it difficult to identify its specific benefits. We found that this characteristic was common among expert systems used as closed-loop controllers, that is, those systems with benefits that cannot be traced to the component level. This observation suggests that expert systems have taken their place among the suite of techniques available to the controls engineer and do not require the special attention sometimes afforded new technologies.

This observation suggests that expert systems have taken their place among the suite of techniques available to the controls engineer and do not require the special attention sometimes afforded new technologies.

Types of Application

Although a wide variety of knowledge-based systems were presented to the panel, the systems generally fell into a few categories: diagnosis and troubleshooting, planning and scheduling, configuration tasks, process monitoring and control, and software engineering.

Diagnosis and Troubleshooting

Diagnostic systems comprise a broad range of applications that deduce faults and suggest corrective actions for a malfunctioning system or process. Numerous examples can be found in physical (electronic, mechanical, hydraulic, and so on) and biological domains as well abstract domains such as organizations and software.

A high proportion of the earliest knowledge systems were diagnostic systems or classification systems. Although the emphasis has now shifted toward planning and scheduling systems, which often produce greater benefits, diagnostic systems still make up over 40 percent of the total number of

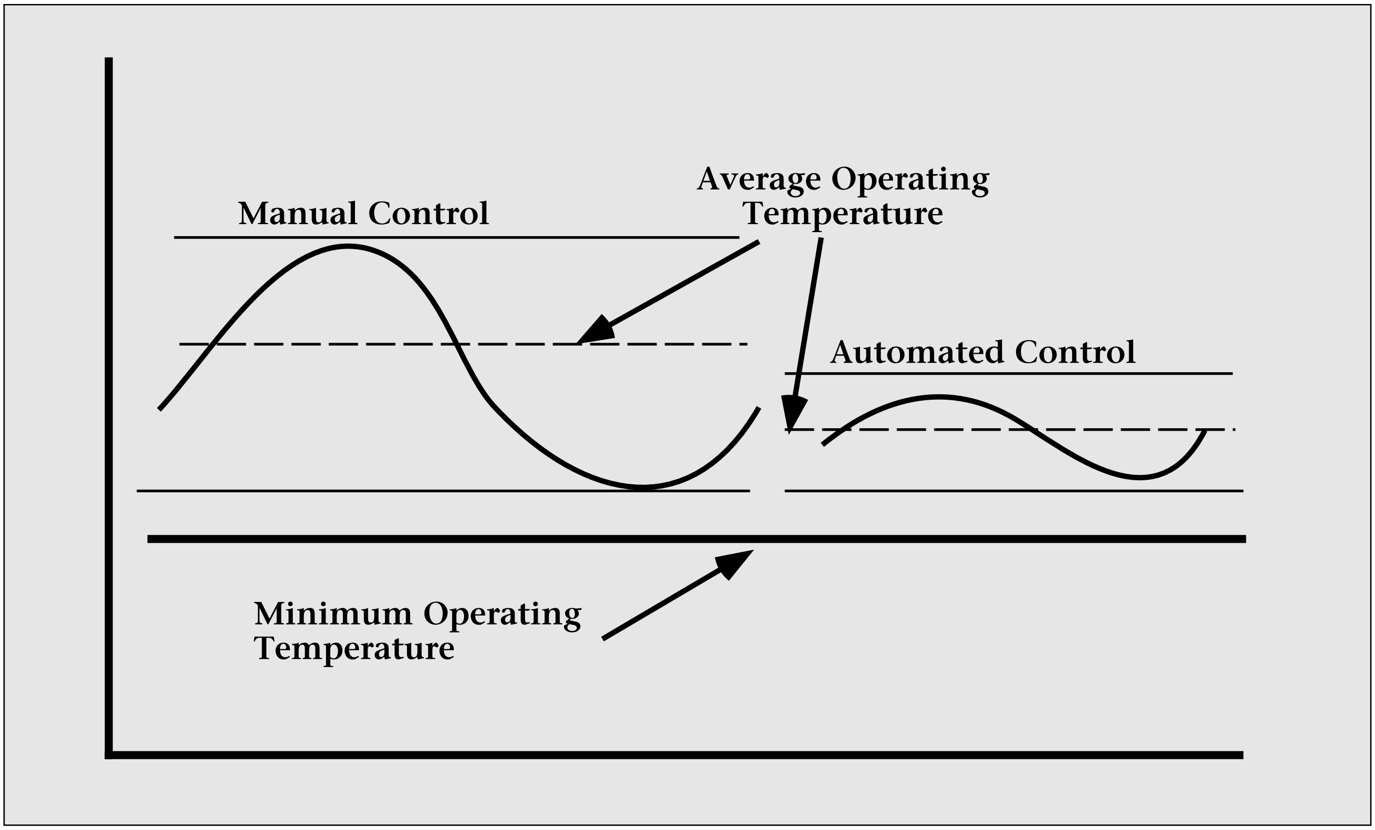

Figure 2. Fuel Cost Savings.

fielded systems (*Nikkei AI* 1992). It is likely that diagnostic systems will continue to represent a substantial proportion of knowledge system applications because tool vendors have produced a number of task-specific shells for diagnostic applications (see Tools and Infrastructure). These tools will make the development task easier and will broaden the base of people capable of developing diagnostic applications.

We learned of major applications in electrical distribution system fault isolation, nuclear power plant diagnosis, telephone cross-bar–switch diagnosis, and subway air conditioning diagnosis. Each of these applications is a large system offering substantial benefits. One application handles 500 different kinds of electromechanical equipment consisting of 25,000 component types.

Japanese steel producers have developed a number of diagnostic expert systems. One important application is the determination of the current state and trends in blast furnaces. Although not a complete diagnostic system, the characterization of state is the first step in a diagnostic system.

Several applications in the domain of computer systems were mentioned by the Japanese computer manufacturers. These applications included determining the best means to recover from system failures and offering advice on software debugging.

In construction, the process of determining the cause of concrete cracking has been automated (Obayashi Corp.). Applications in medical diagnosis and finance were also mentioned but not detailed.

The benefits of diagnostic applications include (1) reduced downtime, (2) safe recovery from failures, and (3) accumulation or preservation of knowledge. In the case of the third point, it was specifically mentioned that the more educated young Japanese do not want to make their careers in certain operational jobs. Capturing the knowledge of current experienced workers is essential to future operations using less educated labor. In one case, the cross-bar–switch diagnostician, the application eliminates the need to train replacement personnel for this obsolete equipment.

Planning and Scheduling

A strong emphasis on this class of systems was apparent at many of the sites we visited. Based on the corporate estimates, somewhere between 30 and 50 percent of the more recently fielded knowledge base systems were for planning or scheduling, with a significant recent trend toward more applications. This trend might be because such systems are used routinely—factory and airline schedules are created every day, week, or month—and success is relatively easy to measure—better schedules (that is, the work in progress time or resources used to accomplish some task is less) or quicker production of schedules. For example, crew-scheduling time at Japan Airlines was reduced from 20 to 15 days; scheduling time for a Toshiba paper mill was reduced from 3 days to 2 hours; a Fujitsu printed circuit board assembly and test planner reduced the scheduling task by a person-year each calendar year; and for an anonymous Hitachi task, scheduling was reduced from 17 hours to a few minutes.

All the systems we were able to examine in detail used straightforward heuristic scheduling methods. In most cases, a complete and correct schedule was good enough; there was little emphasis on optimization of schedules.

Although several sites mentioned a need for reactive rescheduling methods, we did not observe a currently operational system with these capabilities (however, the elevator-group control system described later can be considered an example of a highly reactive scheduling system). Within the United States, iterative improvement methods such as simulated annealing are now being used to effect anytime rescheduling. These methods quickly produce a complete but possibly poor schedule and improve the schedule until available time runs out, always maintaining a correct schedule when interrupted.

Configuration of Manufactured Objects from Subassemblies

We saw little application of expert systems to configuration-type problems such as those that occur in design or manufacturing. The most prominent examples were the systems developed at Sekisui Heim for modular housing, and a system developed at NTT for the design of private telecommunication networks. Fujitsu is planning expert systems for computer-integrated manufacturing but did not elaborate. NEC has investigated rule-based and algorithmic approaches to large-scale integration (LSI) design and has developed EXLOG, a system for synthesizing customized LSI circuits and gate arrays (Iwamoto et al. 1991).

Process Monitoring and Control

The most significant example of a knowledge-based system for control was the one installed in 1987 at the NKK blast furnace, as described previously. Since this time, the steel and construction industries in Japan have been active in developing knowledge-based control systems. These systems are characterized by their high degree of integration with conventional information-processing systems and the widespread use of fuzzy control methods.

Other examples of knowledge-based control systems include one developed by Mitsubishi Electric for control of a group of elevators, another for automatic direction control of a shield tunneling machine, and one for control systems for a swing cable (Obayashi).

Software Engineering

Improving the process of developing software is potentially one of the most highly leveraged applications for new technology. Several companies indicated that 5 to 10 percent of their applications were in the development, testing, or management of software. However, we did not visit any software factories, and we saw only three brief demonstrations of software-related applications. Most research in the software-engineering laboratories focuses on new forms of software as opposed to knowledge-based techniques to support the development of conventional procedural software.

The examples we saw included the generation of procedural programs from a state-event matrix and from a problem model. Both of these applications utilize transformation technology to transform a declarative specification into procedural code. Both operate at the individual program level rather than the system level.

Another example was NEC's application of case-based reasoning to the retrieval of software-quality–improvement ideas. The case base was developed over many years, which, in itself, is a unique contribution. Historically, the case base, which is company confidential, has been published in a book that is updated annually. However, the book has now become too large to be effective, and an electronic case library has been established. The library indexes using 130 attributes, and some clever techniques help minimize the number of index items that the user confronts. We did not view this application as a convincing example of case-based reasoning because there is no repair logic in the application.

None of the applications we saw or heard about represent innovations that go beyond applications in the United States. The use of knowledge-based techniques to support the development of conventional software does not appear to be a priority task in the companies we visited.

Observations and Conclusions

The technology of expert systems has now been mastered by the Japanese. Since the early 1980s when they first entered this field, they have completely caught up with the United States. Their best applications are equal to the best elsewhere in the world. Their use of the technology is widely spread across many business categories.

The number of fielded applications in Japan is somewhere between 1000 and 2000, including personal computer–based applications. The number of U.S. applications is probably several times that of Japan.

Within the computer-manufacturing companies, there is a close coupling between activities in the research laboratories, the system development groups, and the sales departments. The development and sales groups work closely together to develop custom systems for clients, the results of which are fed back to the research lab to provide the requirements on the next generation of expert system tools.

Viewed as a technology (rather than as a business), the field of expert systems is well established in Japan, as it is in the United States. As in the United States, the experimentation phase is over, and the phase of mature applications is in progress. Following a normal learning curve, the number of successful deployments of expert systems has risen sharply, from about 5 percent in the early years to about 75 percent in recent years. Japanese appliers of the technology make eclectic use of AI techniques, most of which originated in the United States or Europe. As in the United States, expert system technology is often a component of a bigger system. The Japanese do not attempt to analyze payoff at the component level but at the system level. Thus, they do not measure the return on investment of these embedded expert systems. However, there are many applications in which the expert system is the main technology.

Viewed as a business, the expert system field did not take off in any exceptional way like in the United States or Europe. Although the overall level of activity is significant and important, there is no evidence of exponential growth. The components of the business consist of expert system tools, consulting, and packaged knowledge systems. Hitachi's expert system business seems the most viable. Other major players, such as Fujitsu and CSK, have not had business success.

Tools and Infrastructure

An expert system tool, or a *shell*, is a software development environment containing basic components of expert systems. Associated with a shell is a prescribed method for building applications by configuring and instantiating these components. Some of the generic components of a shell are the knowledge base, the reasoning engine, the knowledge-acquisition subsystem, the explanation subsystem, and the user interface.

First-generation expert system tools were general purpose and often distinguished by the platform on which they ran (for example, see Harmon 1992a, 1992b). In recent years, new tool types have come on the market that are specialized according to task (for example, diagnostic) and problem-solving approach or solution type (for example, case based). These second-generation tools encode the problem-solving know-how gained by building applications in different areas using the first-generation tools. A problem-specific or task-specific tool contains (1) knowledge representation schemes and reasoning methods found useful for a particular class of applications and (2) a task ontology associated with the problem class.

Table 1 lists most of the commercial tools developed in Japan. They are broadly categorized as general purpose, task specific, solution specific, and development methodology (that is, tools for training those implementing in the methodology for developing expert systems). More general-purpose tools are on the market than the list might indicate. A general-purpose tool such as ES/KERNEL represents a class of tools, and a version of the tool exists for different types of hardware platforms—ES/KERNEL/W for the workstations, ES/KERNEL/H for the mainframes and supercomputers, ES/KERNEL/P for the personal computers, and ES/KERNEL/D for online data processing. The four most popular tools are those developed by domestic computer manufacturers: Hitachi, ES/KERNEL and EUREKA; Fujitsu,

Tools	Company	Remarks
General Purpose		
ESHELL	Fujitsu	Lisp-based (UTLisp), primarily a mainframe tool
ES/KERNEL	Hitachi	In C on UNIX workstations
EXCORE	NEC	
BRAINS	TIS	
XPT-II	CSK	
KBMS	NTT	
ASIREX	Toshiba	Second-generation tool in C on workstations
Task-/Domain-Specific		
APSHELL/DIAG	Fujitsu	Diagnostic
APSHELL/GUIDE	Fujitsu	Consultation
APSHELL/SCHEDES	Fujitsu	Scheduling
EUREKA	Hitachi	Real-time process control
ES/PROMOTE/DIAG	Hitachi	Diagnostic, built on top of ES/KERNEL
ES/PROMOTE/PLAN	Hitachi	Planning and scheduling
HPGS	Hitachi	Scheduling
DIKAST	Toshiba	Diagnostic with knowledge-acquisition facility
PROKAST	Toshiba	Scheduling with knowledge-acquisition facility
DEBUT	Toshiba	Parametric design
PLANBOX	NEC	Planning
FTX	Ishikawajima-Harima	Diagnostic
ESTO	Nippon Steel	Diagnostic
GENZO	Shimadzu	Diagnostic, classification, interpretation
GENZO-II	Shimadzu	Design and planning
GENZO-QAE	Shimadzu	Q/A systems
MEL-DASH	Mitsubishi	Electric network diagnostic
Solution-Specific		
ES/TOOL/RI	Hitachi	Rule induction (under development)
EXCEEDS3	Hitachi	Qualitative reasoning (under development)
Development Methodology		
ES/GUIDE	Hitachi	Text-based explanation of development methodology
ES/SDEM	Fujitsu	
TUPPS-ES	Toshiba	
ES/METHOD	NEC	
SOLOMON	Mitsubishi	
Others		
FORTRAN/KR	Fujitsu	FORTRAN-based, mainframe and workstation tool
YPS/KR	Fujitsu	COBOL-based, mainframe tool
KWSHELL	Fujitsu	Knowledge-acquisition tool

Table 1. Commercial Expert System Building Tools in Japan.

	1987	1988	1989	1990	1991	Total
Scheduling	2	7	15	5	2	31
Control	5	10	7	9	20	51
Diagnosis	5	6	2	15	11	39
Fuzzy	2	3	4	9	12	30
Neuro	0	0	2	6	25	33
Total	14	26	30	44	70	184

Table 2. Expert System Development History of Top-5 Expert System User Companies (R. Mizoguchi, personal communication).

ESHELL; and NEC, EXCORE.

The list of tools in table 1 shows the recent emergence of task-specific or domain-specific tools. The decision to invest in the building of a task-specific tool is based on the demand for application systems in the area. For historical reasons, diagnosis was the most popular application area for the first few years following the introduction of expert systems. Thus, diagnostic problem solving is a well-understood task area. More recently, both Fujitsu and Hitachi claim that scheduling and planning systems are becoming popular, and there are demands for specialized tools in these areas. Toshiba claims that design systems, especially in the area of LSI routing, are also increasingly in demand.

An interesting phenomenon is the popularity of Hitachi's EUREKA, a tool for developing real-time control systems. It outsells domain shells in all other task areas, even though the number of real-time control systems being built is much smaller than other types of applications. Table 2 shows the actual numbers of applications developed by the top five companies using ES technology. The table shows a surge in 1991 in the use of control programs, which goes hand in hand with the increase in the use of neural networks. The popularity of EUREKA parallels the rising popularity in the United States of G2, a real-time control system tool developed by GENSYM. (However, the use of neural networks in control systems is still rare in the United States.)

Hitachi's EUREKA … outsells domain shells in all other task areas, even though the number of real-time control systems being built is much smaller than other types of applications.

Infrastructure

Most of the companies that the JTEC team visited were large enough to have their own research laboratories. Like their counterparts in the United States, industrial laboratories of large companies in Japan engage in basic research and advanced development in a variety of areas. Toshiba, for example, has five research laboratories that conduct long-range (5 to 10 years) research in areas from computer hardware to manufacturing to software. In addition, Toshiba has eight development laboratories that look ahead three to five years. In 1991, with annual net sales of ¥27.3 trillion (US$22.9 billion), it spent 8.2 percent on research, nearly double the amount in 1981.

Observations and Conclusions

In general, Japanese tool vendors are optimistic about expert system technology. Hitachi appears optimistic about expert systems as a business as well. Although AI fever has cooled, there is now a better understanding of the strengths and shortcomings of the technology. There are fewer exploratory users and more users demanding practical systems. There is also a steady increase in business.

The majority of Japanese expert system tools are developed, sold, and applied by computer companies. They have the resources to conduct research, develop new products, and persist in the business.

Because of the close relationship between industrial research, system development, and sales personnel in Japanese companies, solutions to customer problems are identified cooperatively and then quickly find their way into expert system tools.

Many Japanese tools under development are at about the same level of sophistication as American tools. Although many new ideas originate in U.S. research laboratories, the Japanese are quick to develop the technology and transfer it into products.

The predominant application areas have been equipment diagnosis, planning and scheduling, design, and process control. As in the United States, the Japanese companies are developing task-specific or domain-specific tools for diagnostic, planning and scheduling, and real-time control problems. These tools are often combined with knowledge-acquisition aids.

As in the United States, the Japanese are moving toward open, client-server architectures. The impact on the tool business is an increased demand in workstation tools, especially UNIX-based tools written in C. Concurrently, there is a slowdown in the demand for mainframe tools.

All the major computer companies conduct research in knowledge-based systems. Most of the research is in applying or integrating new techniques to customer problems. The industrial research laboratories serve as technology-transfer agents for both imported and internally developed techniques and methodologies. At the same time, as with consumer products, the Japanese companies are spending research money on improving and refining ideas and products.

On the negative side, the Japanese suffer from a proliferation of tools, an image that reflects their computing industry: (1) there are several large computer manufacturers whose hardware products are incompatible, (2) customer loyalty keeps end users from shopping around, (3) customers tend to desire custom systems, and (4) there does not appear to be any movement toward standardization. However, with the move toward open-system architectures, these patterns might significantly be altered, and one or two dominant players might appear. Hitachi, with its UNIX workstation tool, is showing signs of being one of them.

University Research in Advanced Knowledge-Based Systems

Although the best research in Japan compares favorably with that in the United States, the base of fundamental AI research in Japan is far less broad and less deep. A reasonable measure of comparison is the number of publications in the proceedings of the premier conference in the field, the International Joint Conference on Artificial Intelligence (IJCAI). An analysis of the most recent three IJCAI conferences (Australia in 1991; Detroit, Michigan, in 1989; and Italy in 1987) reveals the following results.

There were 37 Japanese single or coauthored publications over this time span, compared to 387 American publications. Of these publications, 18 came from academia (9 from Osaka University; 6 from Kyoto University; and 1 each from Kyushu University, the University of Tokyo, and the Tokyo Institute of Technology). The remaining 19 publications came from government laboratories (5 from ICOT and 4 from ETL) and industry (5 from NTT, 3 from NEC, and 1 each from Hitachi and Toshiba). It should be noted that the industrial papers were heavily clustered in the fields of robotics (particularly machine vision) and natural language (both areas beyond the scope of the knowledge-based systems theme of this report). A complete analysis of the U.S. publications was not undertaken, but roughly 75 percent were from academia, encompassing at least 50 different sites. Over the same time period, Germany produced 50 IJCAI publications; Canada, 48; the United Kingdom, 39; France, 35; and Italy, 23.

While in Japan, the JTEC team visited both Kyoto and Osaka Universities as well as a relatively new center for research at the University of Tokyo. Based on the previous statistics, we believe this sampling provided reasonable insight into the overall Japanese academic AI research establishment.

Research at Kyoto University, under the direction of Toyoaki Nishida, focuses on symbolic reasoning about the behavior of physical devices and processes when formal mathematical models are either unknown or computationally intractable. Nishida has recently moved from Kyoto University to head the Knowledge-Based Systems Group at the Advanced Institute of Science and Technology in Nara, one of two new graduate schools recently created in Japan. It is worth noting that each of these graduate schools (the other is the Japan Advanced Institute of Science and Technology and is located in Hokuriku) will have 20 computer science research groups, with the traditional four professors for each group.

At Osaka University, Riichiro Mizoguchi has created an entrepreneurial American-style AI research laboratory. He supervises eight Ph.D. students, a large number by Japanese standards. He often has foreign visitors and has raised substantial sums of industrial cofunding for his laboratory.

Mizoguchi's laboratory is conducting research in four areas of knowledge-based systems work: (1) the use of deep knowledge in next-generation expert systems; (2) knowledge acquisition (building an interviewing system capable of automatically instantiating an expert system for a particular task with no intervention by a knowledge engineer); (3) large-scale, reusable and shareable knowledge bases; and (4) intelligent tutoring systems.

At the University of Tokyo, the Research Center for Advanced Science and Technology, has five focus areas. One of these areas, knowledge processing and transfer, is under the direction of Setsuo Ohsuga. The lab has 18 graduate students, 5 of whom are non-

Japanese, and 5 research staff members (2 of whom are foreign visiting scholars). Much of the work in this lab is conducted in conjunction with industry consortia. The primary research focus in Ohsuga's lab is on knowledge representation for intelligent computer-aided–design applications across a variety of domains.

In summary, a survey of three years of working papers of the Special Interest Group on Knowledge-Based Systems of the Japan Society for AI shows a wide range of research topics, touching most of the subjects of current interest in the United States. The quality of research at a few top-level universities in Japan is in the same range as at top-level U.S. universities and research institutes. In the remainder of the Japanese university system, the quality of research is not at the same level as at first- or second-tier U.S. research centers. However, the quantity of research (in terms of number of projects and number of publications) is considerably smaller (by nearly an order of magnitude) compared to the United States.

Industrial Research

Many industrial laboratories tend toward advanced development work or are windows on a fast moving world of scientific information. These tendencies are pronounced in Japan. Work in knowledge-based systems in Japanese industrial laboratories appears to be tightly coupled to application or product development. Japanese computer-manufacturing companies and certain high-technology companies carry out some knowledge-based systems research. Other types of companies do virtually none. The JTEC team observed a thin layer of excellent industrial research at Fujitsu, Hitachi, Toshiba, NEC, and NTT. From other publications, we know that there is excellent knowledge-based systems work at IBM Japan and Sony. Perhaps the most extensive and the deepest of the basic research activities at companies was seen at Hitachi's Advanced Research Laboratory (under the direction of Hiroshi Motoda), NEC, and NTT.

National AI Research Projects

The panel visited the Electronic Dictionary Research (EDR) Project; ICOT, which is the laboratory of the Japanese Fifth-Generation Computer Systems (FGCS) Project; and LIFE, which is the Laboratory for International Fuzzy Engineering. The panel also looked into a new national project, called Real-World Computing (RWC), which is a successor to the FGCS project.

EDR

EDR was spun from ICOT (in 1985) with a nine-year charter to develop a large-scale, practical electronic dictionary system that could be used in support of a variety of natural language tasks, for example, translation between English and Japanese, natural language understanding and generation, and speech processing.

EDR is well on the way to completing a set of component dictionaries that collectively form the EDR product. These dictionaries include word dictionaries for English and Japanese; the concept dictionary; co-occurrence dictionaries for English and Japanese; and two bilingual dictionaries, English to Japanese and Japanese to English. Each of these dictionaries is extremely large scale. The concept dictionary, for example, which captures all semantic information in a form that is independent of the surface language, contained approximately 400,000 concepts as of the time of the panel's visit.

In contrast to more theoretically based approaches taken in the United States, the construction of the EDR dictionaries has been a largely pragmatic effort. Tools to automatically index a concept within the semantic network, check it for consistency, identify duplication, and so on, are usually simple ones that can identify potential problems, relying on humans to resolve these problems. Despite the weakness of their tools, there is a sense that closure is being reached; that is, most words and concepts currently encountered are already found in the dictionary.

The JTEC team found EDR's pragmatic approach refreshing and exciting. Although U.S. researchers have spent significantly more time on theoretic formulations, they have not yet succeeded in building any general knowledge base of significant size (the number of named terms in the CYC system is still significantly smaller than the EDR concept dictionary, although the knowledge associated with each term is much deeper). EDR's pragmatic approach (and ability to enlist a large number of lexicographers into a single national project) has allowed it to amass a significant corpus of concepts with significant coverage of the terms of natural language. The organizing ideas of the EDR dictionary are not particularly innovative; they have been in play since Quillian (mid-1960s) and Schank (late 1970s). Although stronger theoretical work and better tools are both necessary and desirable, there is no substitute for breadth of coverage and the hard work necessary to achieve it. EDR has uniquely achieved this breadth. Its work is among the most exciting in Japan (or anywhere else).

A follow-on project, the Knowledge Archives Project, aims to amass large knowledge bases to serve as the "...common base of knowledge for international and interdisciplinary exchange in research and technology communities."[2] This goal will be achieved primarily by (semi)automatically processing textual knowledge sources (that is, documents) into large-scale semantic networks. The knowledge base will support a variety of knowledge storage, retrieval, and transformation tasks. Largeness is seen as the core opportunity and challenge in the effort. CYC is the U.S. project most similar to the one proposed here, although it is important to note the difference in perspective (Lenat and Guha 1990). In CYC, it is the job of human knowledge engineers to develop an ontology and enter the knowledge into it. In the knowledge archives project, the source of knowledge is to be existing textual material, and the ontology should (at least somewhat) emerge from the self-organization of the knowledge base.

The proposal explicitly mentions the intention of collaborating with researchers outside Japan and encouraging the formation of similar efforts in other countries. The proposal has not yet been approved, but its progress should be followed.

ICOT

ICOT was founded as the central research laboratory of the FGCS project in 1982. After reaching the end of its planned 10-year lifetime in 1992, the project was extended for an additional 3 years with a reduced staffing level. (MITI has also launched RWC to succeed the its FGCS project; however, knowledge-based systems are not a part of the RWC concept.)

The FGCS project was motivated by the observation that "...current computers are extremely weak in basic functions for processing speech, text, graphics, picture images, and other nonnumeric data, and for artificial intelligence type processing such as inference, association, and learning" (ICOT 1982). To address these shortcomings, the FGCS project was commissioned to build the prototypes for a new (the fifth) generation of hardware and software. Early ICOT planning documents (ICOT 1982) identify the following requirements: (1) realize basic mechanisms for inference, association, and learning in hardware and make them the core functions of the fifth-generation computers; (2) prepare basic AI software to fully use the stated functions; (3) take advantage of pattern-recognition and AI research achievements and realize man-machine interfaces that are natural to humans; and (4) realize support systems for resolving the software crisis and enhancing software production.

A fifth-generation computer system in this early ICOT vision is distinguished by the centrality of problem solving and inference, knowledge base management, and intelligent interfaces. Such a system would require significant advances in both hardware and software. In practice, the core ICOT efforts focused on parallel symbolic programming (in particular, parallel logic programming) by developing a new language (called KL1) and developing experimental hardware to support the language (the parallel inference machines [PIMs]). Here, we summarize ICOT's accomplishments by answering several questions:

Has ICOT directly accelerated the development of knowledge-based technology in Japan? Based on the various demonstration applications that were developed as part of the FGCS project and the relatively small amount of effort placed on knowledge acquisition and learning, the answer would be no. ICOT made little progress in the application dimension and has had little impact on knowledge-based systems technology. The choice of Prolog and logic programming isolated ICOT from industry. ICOT machines are research-only designs. Their high cost has prevented distribution outside ICOT and consequently has isolated ICOT from the industrial and commercial communities.

Has ICOT indirectly affected the state of knowledge-based technology in Japan? Certainly, the answer is yes. The establishment of a national project with a focus on fifth-generation technology attracted a great deal of attention for AI and knowledge-based technology. Several sites commented on the fact that ICOT had attracted better people to the field and lent an aura of respectability to what had previously been regarded as esoteric.

Is the ICOT work likely to produce a platform that will ultimately accelerate knowledge-based technology in Japan? This question is difficult question to answer.

The FGCS project was motivated by the observation that "...current computers are extremely weak in basic functions for processing speech, text, graphics, picture images, and other nonnumeric data, and for artificial intelligence type processing such as inference, association, and learning"

ICOT's work has built an elegant framework for parallel symbolic computing. Most AI people agree that without parallelism, there will ultimately be a barrier to further progress because of the lack of computing power; however, this barrier does not seem imminent.

Has ICOT's work advanced the state of parallel-processing technology in Japan (or elsewhere)? One can argue both sides of this question. On the positive side, ICOT has developed a parallel symbolic computing language, KL1, and virtually no interesting work has been done elsewhere for expressing parallel symbolic computation. Another positive point is that ICOT will have a test bed of several PIM machines to explore over the next year. No other site has access to several distinct implementations of the same virtual parallel machine. It is not unreasonable to expect significant insights to emerge from this experimentation. On the negative side, the ICOT work has tended to be a world closed in on itself. In both the sequential and the parallel phases of the research, there has been a new language developed that is only available on the ICOT hardware. Furthermore, the ICOT hardware has been experimental and not cost effective. Consequently, the ICOT technology has had little or no impact on, or enrichment from, real-world applications.

LIFE

Founded by MITI in 1989, LIFE began operations in early 1990 under the leadership of Toshiro Terano. It has three main foci to its program: (1) decision support, which includes work in both human-mediated and fully automatic control systems; (2) intelligent robotics, which includes work in speech and image understanding as well as robot planning; and (3) fuzzy computing, which has the aim of producing the computing hardware and software necessary to fully implement fuzzy systems.

Terano and his colleagues view fuzzy control, which is now embedded in many commercial applications, as the first way station on the road to human-friendly systems. The evolution is fuzzy control to fuzzy expert systems to intelligent robots to integrated AI systems. The 1992 LIFE research plan shows work in two main areas: communication and machine intelligence. Under communication, LIFE is conducting research in fuzzy computing (including natural language and linguistic modeling applications) and intelligent interfaces (to both humans and robots). Under machine intelligence, LIFE is pursuing work in fuzzy associative memory (software only), image understanding, and intelligent robots (the scope of which was left vague in our discussions). About 30 full-time researchers are involved in all these efforts.

Observations and Conclusions

EDR will likely produce a practical-scale, machine-usable dictionary for Japanese and English. With several hundred thousand entries in the concept dictionary, the scale of EDR accomplishments is impressive and should be taken as a model for similar research programs elsewhere. A large follow-up project to EDR on large knowledge bases, called the Knowledge Archives Project, might be funded and should be tracked closely.

EDR has not significantly improved the underlying technology for maintaining large knowledge bases or significantly added to our theoretical understanding of knowledge base organization.

Using massive parallelism, ICOT achieved its stated goal of 100 million logical instructions per second (LIPS) theoretical peak performance (actually achieving 150 million LIPS). On the software side, ICOT is one of only a few sites in the world that is studying massively parallel symbolic computing, and its logic programming research is probably the best in the world.

The Fifth-Generation Project also achieved its goal of training a new generation of computer technologists. Moreover, ICOT created the funding and motivation to spur significant interest worldwide in AI, knowledge-based systems, and advanced computing paradigms.

On the negative side, ICOT made little progress in the application dimension and has had little impact on knowledge-based systems technology. The choice of Prolog and logic programming, coupled with high-cost research machines, isolated ICOT from industry. LIFE is now the world leader in applying fuzzy logic concepts to classic AI core problems.

Notes

1. The complete JTEC panel report, "JTEC Panel Report on Knowledge-Based Systems in Japan," edited by Robert Engelmore, is distributed by the National Technical Information Service of the U.S. Department of Commerce, NTIS Report PB93-170124.

2. This quotation appears in an unpublished report

written by the Economic Research Institute in 1992, "A Plan for the Knowledge Archives Project." It is available from the Japan Society for the Promotion of Machine Industry: Kikai-Shinko Building B1, 3-5-8 Shiba-koen, Minato-ku, Tokyo 105, Japan.

References

Harmon, P., ed. 1992a. *Intelligent Software Strategies* 7(7).

Harmon, P., ed. 1992b. *Intelligent Software Strategies* 8(2).

ICOT. 1982. Outline of Research and Development Plans for the Fifth-Generation Computer Systems, ICOT, Tokyo.

Iwamoto, M.; Fujita, S.; Watanabe, M.; Kawarabayshi, M.; and Maeda, N. 1991. Combining Rule-Based and Algorithmic Approaches in LSI Design. In *Proceedings of the World Congress on Expert Systems,* 1302–1309. New York: Pergamon.

Lenat, D. B., and Guha, R. V. 1990. *Building Large Knowledge-Based Systems: Representation and Inference in the* CYC *Project.* Reading, Mass.: Addison-Wesley.

Nikkei AI. 1992. Special Edition: General Review of Expert Systems in Use. *Nikkei AI,* Special Issue, Winter 1992.

Onodera, K., and Mori, A. 1991. Cockpit Crew Scheduling and Supporting. In *Proceedings of the World Congress on Expert Systems,* 1–10. New York: Pergamon.

Edward A. Feigenbaum is codirector of the Stanford University Knowledge Systems Laboratory. He and his laboratory pioneered the invention and development of knowledge-based systems. He is also director of the Software Industry Research Project of Stanford's Computer Industry Study.

Peter Friedland is chief of the Artificial Intelligence Research Branch of Ames Research Center. He is director of a 60-person laboratory conducting AI research and application developments in the context of aerospace domains. Friedland received his Ph.D. from Stanford University in computer science in 1979. He was awarded the 1994 Feigenbaum Medal for Expert System Research and Applications.

Bruce Johnson is a partner in Andersen Consulting in Chicago. He is currently director of research and development strategy and collaboration. He develops long-term visions and the strategies for realizing these visions. Most recently, he has concentrated on healthcare information infrastructure.

H. Penny Nii, a researcher for 20 years in knowledge-based systems, is best known for her work on blackboard system applications and software development tools. Her work was done at the Stanford University Knowledge Systems Laboratory. Her area of expertise recently shifted from the art of computer programming to the art of quilting. She owns and operates the Penny Nii Gallery in Mountain View, California.

Herbert Schorr is the executive director of the Information Sciences Institute of the University of Southern California, a position he assumed in 1988. He is a graduate of the City University of New York and received his Ph.D. in electrical engineering from Princeton University in 1963. He joined IBM after a year at Cambridge University as a postdoctoral fellow and a year at Columbia University as an assistant professor.

Howard Shrobe is a principal research scientist at the Massachusetts Institute of Technology Artificial Intelligence Laboratory. His current research interests include knowledge-based systems for design as well as intelligent aids for broad-scale communication. Formerly, he was technical director at Symbolics Inc., where he was involved in both software and hardware architectures for advanced symbolic computing.

Robert Engelmore is a senior research scientist and executive director of the Heuristic Programming Project in the Computer Science Department at Stanford University. His primary professional interest is in the applications of AI techniques to scientific and industrial domains. He received a Ph.D. in physics from Carnegie Mellon University. Engelmore was editor in chief of *AI Magazine* from 1981 through 1991 and is a fellow of the American Association for Artificial Intelligence.