Artificial Intelligence (I): Into the World

AI has become a hot property in financial circles; but do the promises have anything to do with reality?

For the people who do research in artificial intelligence—or AI as it is commonly known—these are exhilarating times. From the mid-1950's, when their field first emerged from the postwar ferment in information theory, game theory, cybernetics, and electronic computing, until the late 1970's and the microelectronics revolution, AI was an arcane and academic endeavor with only a few hundred serious practitioners in the world. Today it is big-time high technology.

AI software has finally begun to prove itself in the marketplace, especially with computer vision systems, natural language programs, and "expert" systems that give advice like a human specialist. AI itself has become the cornerstone of plans for a so-called fifth generation of intelligent computers due in the 1990's. The field is awash with venture capital. It has strong support from the Department of Defense. It has been targeted for accelerated development by MITI, the Japanese ministry of international trade and industry. It is a top research priority in Europe. And it is the inspiration for work on totally new kinds of computers, organized in ways not dissimilar to the human brain.

So the future seems limitless—from the outside. Back in the laboratories, however, the hoopla has left researchers bemused and concerned. The demand for trained AI programmers has put a severe strain on the community's manpower resources. Worse, the promises seem to be outracing the reality. At the Xerox Palo Alto Research Center, John Seely Brown puts it very simply: "You've got to separate the science from the hype."

The evolution was not really all that sudden, of course. The clubby atmosphere of AI had begun to change well back in the 1970's, as more and more young researchers entered the field and as the increasing sophistication of AI programs began to stimulate interest on the outside. But for most people it did not really hit home until August 1980, when they arrived at Stanford University for the first annual meeting of the American Association for Artificial Intelligence and found that a meeting planned for several hundred participants was registering nearly a thousand—including an unheard-of 40 reporters and an uncounted number of entrepreneurs and venture capitalists.

"It was as if AI were ready to go public, to get out of the labs," says Louis G. Robinson, the association's first executive director. In retrospect the surge of public interest in AI was probably inevitable: the country's love affair with the personal computer was still fresh and new, industrial robots were making headlines, and microelectronics had become the paradigm of high technology and higher profits; now the Stanford meeting was a chance for outsiders to see what all this "artificial intelligence" was about. But for a community more used to computer terminals than to business suits, the experience was a revelation.

"There was a sense that from here forward, things would not be done in the same way anymore," recalls Robinson, who has since left the AAAI to start his own AI newsletter. And it was true. Maybe after the Stanford meeting people were just more aware of it, but the money seemed to start pouring in. Graduate students and their professors started rushing out to form companies by the dozens, companies with names like Teknowledge (Palo Alto, California), and Thinking Machines Corporation (Waltham, Massachusetts). Medical professionals had gotten interested long before—in fact, much of the work on expert systems has been done by a nationwide community of physicians and programmers working through Stanford's SUMEX-AIM network—but now they were joined by the giant computer firms, the aerospace firms, the mining companies, and the oil companies. The Pentagon, always the principal support for AI through its Defense Advanced Research Projects Agency (DARPA), began to expand its efforts. More recently, at least three companies have started commercial manufacture of computers specialized for LISP, the language of choice for AI applications. And the head-hunting at AI conferences has now become so rampant that people are talking about splitting off AI applications into a special conference by itself, just to give the academic researchers some peace and quiet.

The software, meanwhile, has proceeded apace. INTERNIST-1, for example, an experimental expert system de-

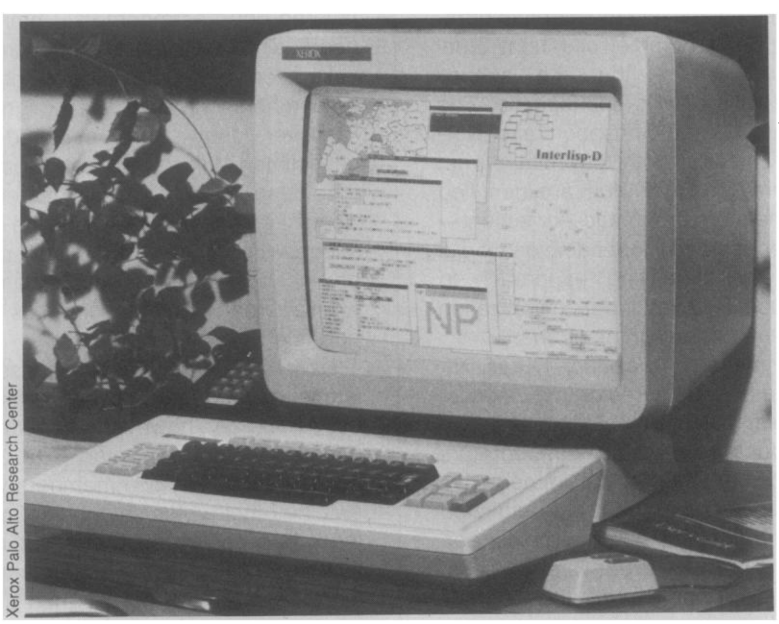

Xerox Palo Alto Research Center

Onscreen "windows" for text and other applications, a handheld "mouse" for the manipulation of what is on the screen—these devices were made famous by Apple Computer's Lisa *and* Macintosh. *But they were invented more than a decade ago by AI researchers.*

veloped at the University of Pittsburgh for diagnosis in internal medicine, was recently tested against cases selected from the *New England Journal of Medicine*. It proved nearly as accurate as the original attending physicians had been (*Science*, 15 April 1983, p. 261). The General Electric Company has gotten very encouraging results from DELTA/CATS-1, an expert system it is developing for the maintenance of diesel locomotives in remote locations; the system is nearly as effective and certainly a lot cheaper than flying the company's one senior (human) expert all over the country.

DEBUGGY, developed at the Xerox Palo Alto Research Center, is in the mainstream of work on computer-aided instruction, one of AI's most active subfields. A far cry from the "programmed" learning of the 1960's and the "computer literacy" efforts of the 1980's, DEBUGGY is a responsive tutor, in this case specialized for subtraction. If a young student is having trouble with subtraction, the program will first isolate his misunderstanding—the "bug" in his mental algorithm—and then will help him correct the problem by making suggestions and responding to his questions. (This is not as easy as it sounds; with even a simple bug the student can get answers that seem like random numbers.)

WEST, also at Xerox, is a computer coach. Students play a simple board game displayed on a video screen ("How the West Was Won"; thus the name of the program). Meanwhile the program watches and builds a statistical model of their behavior. Occasionally, based on its expertise about the game and about good pedagogy, WEST breaks in with tactful suggestions about how to play the game better.

At least three companies are now working on expert advisers for nuclear power plant operators. The idea is that in some future Three Mile Island crisis the operators might be able to talk things over with a calm, unflappable silicon assistant who will monitor the ailing reactor, alert the operators to changing conditions, make suggestions based on its expertise about the plant, and carry out complex orders instantly—without error.

Looking further, pundits have debated long and hard about the ultimate impact of such systems. People have speculated about everything from reading machines for the blind to trucks that drive themselves through the night and unload themselves at their destination. But as AI pioneer Allen Newell points out, they invariably envision intelligent computers as "humans without flaws," fearless and emotionless beings that will simply displace humans from conventional roles. He thinks the reality will be far more complex and subtle—although he makes no pretense to know what that reality will be.

In any case, the continued progress in AI, coupled with parallel advances in microelectronics and computer architectures, has given a lot of people the sense that intelligent computers are finally within our grasp. Thus the impetus for a national, coordinated effort to create them.

That impetus is in no small measure a response to Japan. On 14 April 1982, Tokyo's MITI officially announced its "Fifth Generation" computer project, a 10-year, $1-billion joint venture with the burgeoning Japanese electronics industry. The avowed goal was to leapfrog the current state of the art by producing computers at least three orders of magnitude faster than current machines; expert systems able to tap knowledge bases at least as large and as varied as the *Encyclopaedia Britannica*; a natural language system that can translate from Japanese to English and back; and a high-level interface that can read the written word, recognize images, and converse with human operators in a natural way. [MITI is also sponsoring a parallel effort in very high speed numerical processing, the Super Speed Computer project (*Science*, 6 May 1983, p. 581).]

Now, the irony of all this is that MITI sold its industrial partners on the project largely as a way of catching up with the Americans, who have a commanding lead in computers and commercial AI. But on this side of the Pacific, the announcement came over as a statement of sublime self-confidence from Japan Incorporated after a decade of steady triumphs in autos, steel, and consumer electronics. Western observers went into shock. Expert systems pioneer Edward A. Feigenbaum of Stanford was so agitated that he co-authored a 275-page call to arms entitled, naturally, *The Fifth Generation*. Granted, many experts were skeptical of the Japanese program, especially given Japan's lack of distinction in basic research. The fact was that we had nothing comparable. American and European research was fragmented among dozens of separate institutions, went the argument. And if the Japanese were to achieve even a fraction of their goals, they could gain a decisive competitive advantage in the 1990's.

The response was swift. Britain launched its own version of the Fifth Generation project, the so-called Alvey Programme, while researchers on the Continent started pushing hard for an international program to be called ESPRIT (*Science*, 6 Jan., p. 28).

In this country, meanwhile, 13 companies took advantage of the Reagan Administration's relaxed attitude toward the antitrust laws to form a research consortium known as the Microelectronics and Computer Technology Corporation (MCC). Under MCC president Admiral Bobby R. Inman, former deputy director of the CIA, the group plans to spend some $1 billion over the next 10 years on AI, advanced microelectronics, and advanced computer architectures.

And, in Washington, the hoopla finally gave DARPA the leverage it needed to break loose Pentagon funding for its 10-year, $1-billion "Strategic Computing" program (*Science*, 16 Dec. 1983, p. 1213). Long enthusiastic about advanced computer technology in the face of skepticism from the Defense Department hierarchy, DARPA now has its institutional neck well extended. The plan prom-

"You've got to separate the science from the hype."

Further Reading

The Handbook of Artificial Intelligence (Kaufman, Los Altos, Calif.)
- Volume 1, Avron Barr and Edward A. Feigenbaum, Eds. (1981)
- Volume 2, Avron Barr and Edward A. Feigenbaum, Eds. (1982)
- Volume 3, Paul R. Cohen and Edward A. Feigenbaum, Eds. (1982)

The Fifth Generation, Edward A. Feigenbaum and Pamela McCorduck (Addison-Wesley, Reading, Mass., 1983)

Machines Who Think, Pamela McCorduck (Freeman, San Francisco, 1979)

An Overview of Artificial Intelligence and Robotics, William B. Gevarter
- Volume 1, *Artificial Intelligence*, June 1983, NASA Technical Memorandum 85836
- Volume 2, *Robotics*, March 1982, National Bureau of Standards, NBSIR 82-2479
- Volume 3, *Expert Systems*, May 1982, NBSIR 82-2505
- Volume 4, *Computer Vision*, September 1982, NBSIR 82-2582
- Volume 5, *Natural Language Processing*, April 1983, NASA Technical Memorandum 85635

ises such wonders as an autonomous land reconnaissance vehicle, an expert system for naval battles, and a silicon pilot's assistant—all within 10 years—plus an industrial infrastructure that can both support the Pentagon's needs and develop *non*-military spin-offs.

The American AI community has paid especially close attention to the DARPA program because they rely so heavily on that agency for funding. But feelings are mixed. On the one hand, enthusiasm for this kind of coordinated development program is high. It promises to infuse the community with money, direction, and a powerful new generation of computing tools. "[The computers being planned by DARPA] will be like the difference between a desktop microcomputer and having a Cray-1," notes robotics expert Raj Reddy of Carnegie-Mellon University.

But on the other hand, there is concern that the DARPA plan may be unrealistically ambitious. Worse, DARPA, the MCC, and all the AI commercialization efforts are producing an enormous boost in applications-oriented research without any increase in funding for basic research or manpower development. The AI community is still quite small. Who is going to do the work? DARPA, for example, is explicitly leaving its support for AI's basic research at the current level of $15 million per year; agency managers are reluctant to ask the Pentagon brass for increases in that area at the same time they are launching an expensive new program.

"AI is in chaos," says Roger Schank of Yale University and co-founder of Cognitive Systems, Inc., a firm that produces natural language systems. "It's hard to get good researchers to work on the fundamental problems because the companies are snapping them all up. Theory has stagnated for the moment, and we've lost our momentum."

"I've got no criticism of DARPA spending this kind of money," he adds. "It's critically important that the country mount some kind of concerted effort to develop computer technology. But there's not enough emphasis on the universities, where new researchers are trained. There's a tremendous manpower problem in AI and it's naïve to assume that the people are just going to appear out of nowhere."

Not everyone is so apocalyptic, of course: "The departure of applications-oriented people from the universities to businesses may be quite beneficial to AI," says Nils J. Nilsson, director of the artificial intelligence center at SRI International; "it brings those with applications interest into more intensive confrontation with real problems, and it leaves at the universities a higher concentration of people who are mainly interested in developing the basic science of AI."

Yet nearly everyone in the community does have a sense that the public's expectations for AI have gotten dangerously overheated, with the concomitant risk of disappointment and backlash. From the laboratory, the Fifth Generation/MCC/DARPA cycle looks like a self-exciting system, fueled by media hype. Even the most enthusiastic researchers will admit that nothing in the science of AI has really changed in the last 5 years. Steady progress, yes. But AI has gone commercial because this is the age of venture capital, not because of any fundamental breakthroughs. Patrick Winston, head of the AI laboratory at the Massachusetts Institute of Technology, opens his talks to lay audiences by saying, "If you're a skeptic I want to make you a believer—and if you're a believer, I want to make you a skeptic."

> "AI is in chaos. It's hard to get good researchers to work on the fundamental problems because the companies are snapping them all up."

For an outsider, probably the most striking thing about AI is the way it violates the common notion of what a computer is. Instead of crunching numbers, an AI program uses the computer as a machine to manipulate symbols; instead of following a rigid and precisely defined algorithm, it picks its own way through a problem according to a store of data, facts, and heuristic rules of thumb about the world.

Indeed, it is arguably the most important insight of AI's first two and a half decades that machines can behave intelligently using just two basic ingredients: *search* and *knowledge*. The paradigm is a chess program. At each step the program has to search through all the moves available to it to find a satisfactory one; but because there are some 10^{120} possible sequences in a chess game, the program would be paralyzed unless it had a few rules of thumb to narrow that search to manageable proportions.

Broadly speaking, AI deals with two kinds of knowledge. Factual knowledge, or "book learning," might be represented in the computer as a network of associations: TWEETY is a BIRD is a VERTEBRATE is an ANIMAL, and so forth. Heuristic knowledge, the intuitive rules of thumb derived from experience or passed down from master to apprentice, might be encoded as a maze of logical propositions: IF this condition holds, THEN do that.

To get a feel for the scope and limitations of current AI programs, consider that a human expert—say a chess master—has at his command the equivalent of 50,000 IF-THEN statements. A modern expert system contains at most a few thousand; even the best is still an *idiot savant*.

INTERNIST-1, for example, knew a lot of internal medicine. It understood nothing about physiology or anatomy. Programs such as its successor CADEUCEUS, which have deeper knowledge and which can begin to reason from first principles, are still very much on the forefront of research.

By the same token, existing systems are very narrow, in part because of hardware constraints on computer memory and processing power. So far the programs have been successful only in well-defined and self-contained domains. (In fairness, of course, the same thing could be said of human experts: a lawyer may well be a klutz at auto repair.)

More important still is the fact that none of the existing expert systems can learn in any real sense. The biggest bottleneck in the creation of a new system is the laborious back and forth between the human experts and the programmer as they discover new rules and refine the old ones. Programs that *can* learn are again on the forefront of research.

But most important of all—in every application of AI, not just expert systems—is the lack of anything that might be called common sense (*Science*, 24 Sept. 1982, p. 1237).

"What makes common sense reasoning so difficult is that you need to know so many facts about the world," says Stanford's Feigenbaum. "How many facts? A million? Some people are trying to codify parts of common sense into a qualitative physics or a qualitative psychology. But who's going to take the time to codify all that? Not the companies. What it's going to take is hoards of graduate students working for a generation. That's why the dream of AI since 1955 or 1956 has been to write a program that can learn from experience. That's the right approach—if only we knew how to do it."

Indeed, if two and a half decades of AI research has done nothing else, it has given researchers a sense of awe in the face of the ordinary. Computers seem to have an easy enough time imitating "advanced" human intelligence—systems for playing chess and proving mathematical theorems were among the first AI programs ever written—but they have a terrible time recognizing a human face or understanding a nursery rhyme. The robot has not been built that can walk across a hillside. "We shouldn't be so intimidated by our Beethovens and our Einsteins," says AI pioneer Marvin Minsky of the Massachusetts Institute of Technology. "We're simply so accustomed to the marvels of everyday thought that we never wonder about it."

In part, the mastery of these everyday miracles may just involve the prosaic matter of computing speed, especially in such fields as vision and natural language understanding. A neuron is very slow compared to a microchip, but the brain makes millions or billions of neuronal calculations simultaneously and in parallel; our current generation of serial, one-step-at-a-time computers are hopelessly outclassed. Some of the most intriguing AI research involves the efforts by many groups to build machines that can do this kind of parallel processing on a suitably massive scale—and to figure out how to program these machines sensibly once they are built.

But many AI researchers, Schank and Minsky among them, think that fundamentally new approaches are needed. Whatever is going on within our skulls when we learn something or when we figure something out, whatever is involved in recognition and memory, it is not a series of neuronal IF-THEN statements. "The thing is, AI is very hard," says Schank. "What is the nature of knowledge? How do you abstract from existing knowledge to more general rules? How do you modify the knowledge when you fail? Are there principles of problem-solving that are independent of domain? How do goals and plans relate to understanding?"

"The computer is a way of testing our ideas," he points out. "But first, we need to understand what we're supposed to be building models of."

—M. MITCHELL WALDROP

This is the first in a series of articles on artificial intelligence research. Subsequent articles will deal with such major areas of application as expert systems, machine learning, natural language understanding, and computer vision.

Fertility Hormones Cloned

A group of researchers at Integrated Genetics, a biotechnology firm in Framingham, Massachusetts, has succeeded using recombinant DNA technology to produce two human fertility hormones, human chorionic gonadotropin (hCG) and human luteinizing hormone (hLH). This is one of the first reports of investigators using recombinant DNA technology to produce molecules that are a combination of proteins and carbohydrates in mammalian cells, according to molecular biologist Leroy Hood of the California Institute of Technology. For that reason, says Hood, "I think it's interesting."

The two fertility hormones have similar structures, each consisting of two polypeptide chains that are put together inside cells and "processed." A section at one end of each chain is a marker that guides the chain to the cell's secretory apparatus and is cleaved once the chain gets there. Before the hormones are secreted from the cell, sugar molecules are added to them. The hormone hCG, for example, is 30 percent sugar by mass. If sugars are not added to these hormones, the hormones are biologically inactive.

Bacteria, which molecular biologists usually use as protein factories, cannot carry out this type of processing. Although they can express added mammalian genes, they do not add sugars to the molecules and they do not excrete them. Thus molecular biologists believe that the only way to produce molecules as complex as the fertility hormones is to make them in mammalian cells, using standard methods of genetic engineering. David Housman, a founder of Integrated Genetics and a faculty member at Massachusetts Institute of Technology, used mouse cells to make hCG and hLH, infecting them with a bovine papilloma virus, which inserts itself in the chromosomes of the cells. To the virus, he and his associates added the fertility hormone genes and a mouse metallothionein gene containing control regions that promote gene transcription. These are well-known methods, although, says Housman, to actually make the methods work was a "nontrivial achievement."

The major problem with this method is that the engineered DNA is unstable—the genes tend to rearrange themselves. If this happens, the hormone genes may not be expressed. "We had to be very careful and very persistent to avoid rearrangements," Housman says. "We had to be sure we picked clones that were stable."

Judith Vaitukaitis, an endocrinologist and fertility specialist at Boston City Hospital, has tested the biological activity of the fertility hormones produced by the Integrated Genetics group. "They're quite good," she says. She thinks that these hormones will be clinically useful in the treatment of infertility because they can induce both ovulation and sperm production. Although hCG and hLH are now available for infertility treatment, the hormones are extracted from pituitaries, urine, or placentas and so are not completely pure. Vaitukaitis estimates that there is between 1 and 5 percent cross-contamination with other hormones, which can complicate treatment and clinical research.

The pure fertility hormones also should be of interest to basic research. Robert Canfield of Columbia University's College of Physicians and Surgeons says that, to his mind, one of the more interesting prospects will be to modify the genes at the sites where the sugars attach in order to study how the sugars relate to structure and function. Irving Boimer of Washington University in St. Louis says that he and others would also like to use the cloned hCG to determine the three-dimensional structure of the molecule. "You can't look at the three-dimensional structure of hCG now because there's not enough of it around," he says. Since the fertility hormones are typical of other glycosylated polypeptide hormones, researchers hope that by learning about them they will learn about other such hormones.

In any event, the Integrated Genetics group has shown the feasibility of cloning conjugated molecules in mammalian cells. "It is certainly one very smart approach—no question about it," says John Pierce of the University of California at Los Angeles. "I think it's the way to go."—GINA KOLATA