Artificial Intelligence and the Macintosh

by Dong H. Kim

1. Why is Artificial Intelligence Important?

In January, 1988, the U. S. Commission on Integrated Long-Term Strategy, whose prestigious members included Henry Kissinger and Zbigniew Brzezinski, forecasted that there would be"dramatic developments in military technology" over the next twenty years. The commission's report concluded that "new technology will produce weapons with greater precision, range, and destructiveness" and that "future wars in space will be more intense and rapid, requiring entirely new modes of operation." Such new modes of operation are surely going to be based on various monitoring and guidance systems emerging out of the advanced research projects in artificial intelligence. It is no wonder, therefore, that the U. S. Department of Defense recently designated artificial intelligence (AI) as one of the most crucial technologies to be developed between now and the end of the 20th century.

The Defense Advanced Research Projects Agency (DARPA) has already sponsored a $300 million effort that is dedicated to adapting artificial intelligence to high-speed parallel processors. More recently the same agency has proposed an eight-year $390 million research project in neural network area. DARPA's Deputy Director of the Tactical Technology Office, Jaspo Lupo, went so far as to describe this new project to be more important than that leading to the atomic bomb.

Technological revolutions have often contributed to shifts in national prosperity and geopolitical influence in the world, as was recently documented in *The Rise and Fall of the Great Empires* by Paul Kennedy. Information revolutions driving current economic and military changes are no exceptions. They are likely to prove as fundamental as any past industrial revolution in human history. In turn, AI technology is profoundly changing almost every aspect of today's information technology, affecting global competition, both commercial and military.

In a growing number of industries, computer technology has become a major determinant of competitive advantage. The same is true of weapons industry worldwide. Military operations have come to rely heavily on computer infrastructure. And computerization is as important to the design, production, and maintenance of weapon systems as it is to commercial industry. It is not so much nuclear technology itself but rather the allied computer and AI technologies required for guidance, sensing, and control that are becoming the critical ingredient in the global strategic balance. The terrain-searching guidance system of American cruise missiles has become a vital component of our nuclear strategy and constitutes a prime example of a weapon system in which an AI-based guidance system plays a crucial role.

Cognizant of this emerging strategic situation, Japan has launched its own ambitious artificial intelligence projects, followed closely by the British and the French. In other words, the importance of artificial intelligence is widely recognized mainly because of its unprecedented potential to transform the balance of commercial and military power in the world through the creation of new machines and weapons that can do things never done before.

Today's construction workers move more dirt and build faster not because they are necessarily stronger or have bigger muscles than before but because they have better tools. Likewise, our engineers are designing high-definition television sets and stealth bombers not because they are getting smarter or have bigger brains than before but because they have intelligent software in their computers. Any company endowed with better computer technology and AI-based tools will enjoy advantage because it can develop new and better software ahead of competitors. The faster a company or a nation, for that matter, can develop new and better software and manufacture new and better models of commercial products, the easier it will be to get ahead and stay ahead.

This brings up the important topic of productivity. New technologies have generated more than half of the increases in industrial productivity throughout history. Further advances in computer research and technology are expected to be the main source of productivity gains in the next few decades. The application of artificial intelligence in design, information management, production control will play a major role in achieving such productivity gains.

Annual American gain in productivity declined from 3 per cent in the 1950s to 1 per cent by the early 1980s. In the last decade, the Japanese share of the world semiconductor market has nearly doubled, to 49 per cent, while U. S. share has declined from 55 per cent to 39 per cent. Japan now has cornered nearly 90 per cent of the world market for one mega-bit DRAMS. To cite just one more alarming statistics, between 1955 and 1985, the American share of worldwide machine-tool production declined from 40 per cent to 12 per cent.

If these trends continue, Japan will eliminate America's technological leadership before long. Unless America can boost its productivity growth, it will face a host of new dangers, ranging from spending constraints and falling living standards to continuing dependence on foreign credit and social unrest. The U. S. will have to tailor its economic policies to accommodate foreign creditors whose capital it cannot do without. The United States is now facing a challenge that is basic to its economic prosperity and its status as a world power is at stake.

A critical issue, therefore, is whether the United States can maintain its superiority in AI research and computer technology, thus making sure that its share of gains in future productivity will increase. Artificial intelligence, of which our topic is a corollary, therefore, will be decisive in shaping your future and mine, and its importance as thus stated needs no further emphasis.

Well, this should be enough of an introduction to make you realize how important AI is and make you thirsty for more information. My main purpose this afternoon is to focus our attention on the current status of AI and to assess its prospect. I will then discuss what we can expect of the Macintosh in light of our review and discussion. This first half of my paper will conclude with a brief study of the educational and international implications of artificial intelligence.

In the second part I propose to bring to your attention three case studies of AI software available for the Macintosh today--the first is a brand new and exciting GO program developed by Anders Kierulf at the University of North Carolina called Go Explorer and the other two are AI programs that many of you are already familiar with, namely Instant Expert Plus by Human Intellect Company in California and Cognitron by Cognitive Software of Indiana. Compared to more than thirty development shells available for the IBM PCs, you may think that what is available for the Macintosh is really a slim picking. But equipped with HyperCard and superior AI shells, the Macintosh will have more than enough to catch up and surpass IBM PCs in the coming years.

By the way, some of you may be interested to know that the Chinese GO Association enlisted one of Taiwan's largest computer companies to put up $2 million in trust as the prize money for the computer GO games. The top standing prize is $1 million to be paid to the first computer GO game that can defeat the junior champion of Taiwan. The prize is good for 15 years, so you have plenty of time to go for the prize. Our Explorer Go program, which I will demonstrate later, ranked second in last year's tournament. So those of you who are into game-designing but with little interest in AI should stick around; you will see some exciting GO moves,

Although chess has often been described as the "drosophila" (fruit fly) of artificial intelligence, I believe that the game of GO has a better chance of making a broader contribution to research in AI. Goals, plans, patterns, search, coordination, and perception are important aspects of GO strategy as in various branches of AI.

2. The Popularity of Artificial Intelligence and Expert System

While building computer systems that can exhibit expertise in such areas as medical diagnosis and oil exploration, specialists in artificial intelligence have gained a host of insight towards solving the difficult problems of incorporating specialized knowledge in expert systems. Commercial programs in AI are proliferating very fast. Stripped of the hype that usually surrounds a new technology, AI industry is slowly winning back skeptical consumers with scores of successful programs.

Almost everyone with a doctorate in LISP from MIT, Carnegie Mellon or Stanford seemed to create a new AI company in the late 70s and early 80s. And, as Harvey Newquest pointed out, "There was enough venture capital floating around to fund everyone with a business plan longer than an index card."   Nonetheless, in the last few years more solid AI companies have begun to bridge the gap between overblown promises and actual performances. Consequently, business interest in AI has been rising sharply. After treating AI as if it were a kind of cult magic for many years, corporate America is beginning to acknowledge the utility and benefit of using AI in business.

Many of you may not realize that every time you open a can of Campbell Soup, you are benefiting from the innovative use of expert systems. The canning industry uses huge soup sterilizers called "retort cookers," over 70 feet tall, that can process about 55,000 to 65,000 cans at a time. Campbell operates 130 sterilizers at 8 different cites in the U. S. and in England. If a chain breaks or an obscure belt sticks, it takes a while to figure out what went wrong. The plant staff can handle routine maintenance problems, but the tough ones are passed to Mr. Aldo Cimino, who has worked with the sterilizers for more than 15 years and knows more about them than anyone else.

Aldo usually solves troublesome problems by telephone, but has to fly out to tackle really tough ones from time to time. Whenever that happens, needless to say, Campbell soup waits, losing production and profit. As Aldo was about to retire, Texas Instrument was called in to capture his expertise into an expert system. The result has been so good that the company is currently using expert systems in other areas such as realtime control, scheduling, marketing, and financial applications.

Did you know that NASA is also using an expert system that assesses the risks of thunderstorms near the shuttle launch area. The U. S. Department of Agriculture has a cotton management expert system called Comax. IBM, too, is using over 170 internal expert systems, one of which is employed in diskstorage diagnostics. Having been developed in only 8 months, it is saving Big Blue more than $5 million per year. I sometimes wonder if Apple is using any expert system. If so, how many and what kinds? Perhaps someone from Apple can tell us before this MacHack Conference is over.

One of the most mature and widely used expert systems currently operating on a commercial basis has been developed jointly by an AI group at Carnegie-Mellon and Digital Equipment Corporation. Endowed with more than 6000 rules, this forward-chaining production system written in OPS5 is called XCON and has successfully processed more than 90,000 orders. It can configure all VAX family systems, taking over the tedious jobs previously performed by technical editors, who examined a customer's purchase order to determine what computer components are needed to be substituted or added to make the order consistent and complete. For each order XCON determines necessary modifications, produces detailed diagrams showing the spatial and logical relationships between the hundreds of components, and handles other jobs usually relegated to skilled technical editors. Although it takes a highly-trained, experienced editor as long as 20 or more minutes to finish the job, XCON can typically do it in less than 2.5 minute, including printing!

As computers become less expensive and more intelligent, they are taking on more and more important roles in many fields, ranging from office automation to world-wide integrated banking, communication, and trading systems. Within the computer industry, which is relatively new, the field of artificial intelligence represents even a newer technology. As in all new technologies, adopting ideas and techniques of artificial intelligence will meet resistance because they can be risky as well as productive. Thus, the best way to prove to skeptical managers that AI is worth investigating is to point to the proven results of other successful AI projects such as XCON or Campbell's. Equally successful stories can be told about various systems developed by such major corporations as the Ford Motor Company, General Motors, Arthur Anderson, Corning Glass, Texaco, British Petroleum, Westinghouse, Honeywell.

As to the expert systems under current development, no one really knows how many there are. The National Archives is working on a system designed to handle the routine queries of public researchers, thus relieving its archivists for more creative jobs. A Manhattan law firm is working on another that will facilitate the retrieval of legal briefs and other text materials through apprehension of concepts contained in those documents. The Navy is developing a system to support the decision-making of budget analysts. According to preliminary estimates, with this expert system, 2 to 3 hours of analysis can be finished in less than 15 minutes. DARPA has joined the Air Force Avionics Laboratory in evaluating expert systems that help combat-aircraft pilots cope with emergency. Lastly, but not least, IRS is working on an expert system to help decide whose tax returns to audit.

How many of you remember who won the 1986 Indianapolis 500? Bobby Rahal's March-Cosworth race car won the race. But at the time no one knew that he did so with the help of a newly developed expert system. Strategically located in various parts of his car were sensors that recorded the car's engine rpms, suspension activities, ride heights, gear-shifting, and the aerodynamic pressures. These data were analyzed by the expert system, whose reports were used by the team's mechanics in fine-tuning and adjusting the car.

Indeed, new companies and products in AI are springing up very fast. According to The Source Book of Artificial Intelligence, the number of AI companies has been increasing at the rate of 40% per quarter in the last two years. Current prediction is that AI industry will enjoy a compounded growth rate higher than 30% per year through the mid-90s. It looks as if AI is going to replace desktop publishing as the explosive growth market of the 90s.

3. The Current Status of Artificial Intelligence and its Prospect

Dean Acheson, who was our distinguished secretary of State under President Truman, once described the most exasperating question that can be asked of an American diplomat to be: "What is the foreign policy of the United States?" Similarly, we may characterize the most difficult question that we can ask of any AI researcher to be: What is the precise nature, scope, and purpose of artificial intelligence?" For there is little agreement regarding what artificial intelligence is. One joke making the round in the industry goes like this: "Like beauty, artificial intelligence is in the eyes of the beholder; to academics, it is anything impossible; to the military, it is anything invincible, and to the sales manager of software companies, it is anything improved since the last shipment." At times things in AI can be so chaotic, crash-prone, and unsettling that someone recently posted the following at the end of his UseNet posting to comp.ai SIG: "One year that I have spent in artificial intelligence is enough to make me believe in God."

Unlike most subjects that we study in college, computer science and artificial intelligence in particular do not have a long history. Most of the technical ideas and theories undergirding AI are usually less than 50 years old. It was in 1956 that John McCarthy with the help of his friend at MIT, Marvin Minsky, organized the Dartmouth Conference consisting of "a two-month, ten-men study of artificial intelligence." This was the first known use of 'artificial intelligence' as a new field. According the same announcement, the study was to be based on "the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."

This last statement implies that the main goal of AI is the simulation of human intelligence and learning. Another popular definition is that "The central goals of AI are to make computers more useful and to understand the principles that make intelligence possible." Let's start our inquiry about AI with such broad definitions and assume that we want to make computers useful by making them more intelligent. But how are we to build an intelligent computer when we lack a full understanding of human intelligence? There are so many things that cognitive science has yet to discover about animal and human intelligence and learning. For example, we still do not fully understand how babies learn to speak. We still do not know the interconnection pattern of the brain and how many different kinds of neurons there are in the brain. How, then, can we impart the capacity to learn to our computers when we do not understand or explain how we learn many things?

Optimists have argued against this gloomy view by pointing out that there are many things that we have been able to reproduce without fully understanding their scientific nature or structure. Didn't the caveman build fire without knowing the chemical nature of fire? Isn't it possible to build a clock without understanding any of the physical principles? But learning to build a fire is quite different from simulating intelligence in a modern computer. One cannot simply open a human brain and try to rebuild or simulate it in the same way that one builds a clock by inspecting another clock.

For one thing, every normal human brain has more than 10,000 trillion bytes of long-term memory and about the same for the short term memory. The largest su percomputer systems to be developed in the next decade will not exceed one trillion operations per second, which is about one one-millionth of the computation rate of our brain. Our largest magnetic storage disks hold about one billion bytes of information, which is roughly one ten-millionth of the storage capacity that we have in our brain. In order to simulate the human brain, neuron by neuron, we would need over one million trillion operations per second in computer speed! I am not making these numbers up; I will be glad to give you references.

To this day, AI has had no shortage of enemies and severe critics, who always seem to enjoy gleefully quoting Herbert Simon's 1965 prediction that "machines will be capable, within 20 years, of doing any work that a man can do." Some researchers are now saying that it may take five hundred years to achieve this goal. In defense of Simon, however, we should remember what Jules Verne noted at the turn of the century: "Exaggerated expectations drive scientific progress."

Recently, one critic of AI characterized it as nothing but a "romantic scientific fiction." Someone else has pointed out that most industrial AI applications are using techniques more than 10 years old and that what has made AI suddenly popular is not any special advance in AI research as such but the availability of cheap hardware. In the meantime, some people seem to make a career out of attacking AI and showing what computers cannot do, as is the case with Herbert Dreyfus. Having written his well-known attack on AI in his 1965 Rand report, "Alchemy and Artificial Intelligence," he has authored a more extensive attack on AI in his new book, "Mind Over Machine." His conclusion is that machine will never come close to humans in performing cognitive tasks requiring intuition and holistic thinking. It is becoming something of a family affair now because his younger brother has become the co-author of the book.

The wild optimistic expectations of the early years were tempered by many difficulties and road-blocks. In a sense every major AI project has been a failure at modeling human intelligence. Each successful program has shown what else we must accomplish further. One may even go so far as to say that AI research has been a steady but plodding study of the impossible. The simple-minded attempt at using computers to translate foreign languages ended in disasters. For example, when the sentence, "The spirit is willing but the flesh is weak", was translated into Russian and back into English, it came out as "The vodka is strong but the meat is rotten."

As AI researchers ask questions about language, intelligence, and understanding, they have made us realize how much we do not know. We have also come to appreciate all the more our common-sense knowledge and capabilities. The paradoxical result of AI research is that computers can easily do things that are hard for us, while we can easily do things that are hard for computers. We are poor at remembering details, but we are superb at using over-all rules of behavior and adopting them to new situations as they arise. Exactly the opposite are true of the computers.

The only model we have of a working intelligent system which uses large amount of knowledge is the human. And the hard-wiring of the human brain is the end-product of a billion years of evolutionary adaptation to changing environments. We still do not know how memories have been combined with sensory signals to produce behaviors that enhance the survival of the organism. Right now we do not know enough about the details of how the human mind works to say whether we will ever understand them fully and completely. Therefore, it is risky to predict that we will never be able to build intelligent and creative computers modeled after humans. Neuroscience is bound to come to the aid of AI researchers in the coming years.

As someone recently pointed out, we seem to be going through the Tycho Brahe stage in artificial intelligence. Tycho Brahe is the astronomer who spent forty years collecting data on planetary movements. Researchers in cognitive science, of which AI is a branch, have been accumulating vast amounts of data about how our mind works. One day a cognitive Kepler, Galileo, and Newton may come along and give us a complete understanding of human mind and shed a new light on many remaining problems of AI. Going from the insight to the data is simple; going from the data to the insight calls for great flash of insight that can come only from genius. In the meantime, a series of sustained, cooperative inquiries in all branches of AI must be continued. That is what Newton meant by his remark: "If I have seen further, it is because I have stood on the shoulders of giants."

Only time will tell who will turn out to be giants and midgets among the researchers laboring in AI today. Now that we suffer no more from an inordinate lack of computer memory or other hardware facilities, the only bottleneck blocking progress in AI is going to be the lack of insight and cooperative inquiry based on hard work. And we all stand in need of intellectual breakthroughs, one of which may come from reexamining the standard AI paradigm based on specifying the context, describing the logic of the desired outcome, and trying to achieve it by adopting various heuristics. Approaches under this paradigm will succeed only if the designer has already analyzed the class of problems to be solved and is able to represent this class and the problem-solving heuristics in a suitable programming language. On the other hand, the designers of a multi-layer perceptrons need no precise logical descriptions, only informal understanding of the complexities of the desired behavior sufficient to construct the overall architecture of an appropriate net.

In this regard we should from time to time remind ourselves of the following epigram from John Milton: "The light which we have gained, was given us, not to be ever staring on, but by it to discover onward things more remote from our knowledge." If you study many current variants of back-propagation, Hopfield nets, Boltzmann machines, master-slave nets, plus other multi-layer perceptrons, just to mention a few examples, you will realize that we may be living through a period of really exciting time somewhat comparable to the 1920s and 1930s in nuclear physics. We are indeed pushing the frontiers of our knowledge with regard to our brain.

4. The Place of the Macintosh in Artificial Intelligence

Needless to say in front of this audience here, the very idea of using the desktop metaphor, windows, menus, and icons came from the insight gained by AI people at XEROX's PARC. Ever since, the Macintosh family of computers has been at the cutting edge of the computer industry, partly because of its user friendly interface. Born of the AI research, the Macintosh was indeed a revolutionary milestone. Undoubtedly, the computer historians would say that the Mac constituted a major breakthrough. Recently, Avron Barr, a well-known lecturer and author in AI, went so far to call the Macintosh "An AI machine."

As you know, more and more people are realizing that the Macintosh can be used with little training. If commands are too arcane or complicated to remember between work sessions, the PCs are useless to many people. It is no wonder that, even after five years, the Mac remains an ideal working environment for many users, including programmers.

This fact was well illustrated by a recent UseNet posting by David Oster, who uploaded the following to the comp.sys.mac SIG: "Show me a UNIX or a PC that has anything like the Mac's multi-videocard, multi-color, multi-font, multi-size, multi-style, multi-lingual cut/copy/paste scrolling, mouseable text editor in the operating system. But until you show me one, I 'm going to keep on programming the Mac."

Most of you know that if you buy a new video board for an IBM PC, you have to buy the updates of all your software. In Amiga, as you know, the whole operating system is tied to just one size of display. While other operating systems must start from scratch for a French or German version, we all know that the Macintosh has been multi-lingual all along. Even the latest computers just announced by the Sun, SPARCStation 1 and Sun-3/80, do not yet offer the standard user interface that remain the same across applications. As to the current woes faced by the new IBM PCs, I will not try to enumerate even the most obvious ones. But please allow me to repeat just one telling joke making the round not long ago on the UseNet: "PS/2 is yesterday's hardware today. OS/2 is yesterday's software tomorrow."

When Mikhail Gorbachev returned to Moscow after his first state visit to the United States a couple of years ago, he was asked by someone: "What or who impressed you most out of all the things and people that you saw and met in the U. S. A.?" Do you know what was Mr. Gorbachev's simple but very interesting reply? Believe it or not, his one-word answer was: "Macintosh!"

However, it is too bad that, unlike Mr. Gorbachev, our own government computer buyers remain either ignorant of, or indifferent to, the Mac's deservedly legendary ease of use and training, let alone its amazing capabilities. For example, the U. S. Navy Atlantic Fleet Command in its recent computer purchasing guideline went so far as to require 80286 CPUs and the MS-DOS. Apple would have been disqualified from the bidding from the beginning. Needless to say, Apple protested against this outrageous policy and won the case. Nevertheless, there are still many misguided and ill-informed MIS people within and without the government who continue to regard the Mac as a "crippled, slow, and squinty toy".

Our federal office workers have gone to unusual lengths to obtain the Macintosh, often bringing to the office the Macs purchased for home use. Many users in Federal agencies, ranging from NASA and Budget Bureau to Geological Survey and National Park Service, have been willing to go to such lengths because the Mac provides features they just can't get with IBM PCs, clones or other systems. Chief among these are the Mac's ease of use, its clear and simple graphics and its ability to let users cut and paste information from one application to another. The Mac's ability to manipulate copies for desktop publishing has become legendary among office users everywhere. Apple's software allows integration of text, graphics, sound, music, voice and animation -- all pluses in training applications. In short, the users do not have to become involved in a wrestling match with the software, as is often the case with the other PCs. As someone said, most people would rather have their teeth pulled or drilled than be forced to plod through a tall stack of computer manuals!

It is regrettable that the machine that has been at the cutting edge of computer industry has not been at the forefront of AI research. For various reasons, which I need not discuss here, Apple's entry into AI industry has been very late. The company made its first appearance at the annual meeting of AAAI (American Association for Artificial Intelligence) only last year.

Does this mean that users in AI will have to fight for the Mac in the way that Federal workers have? Some of you may be asking: "Is it too late for Apple to make a difference in AI?" PC users already have windows, and they have just gotten the Presentation Manager as well as a multitasking operating system. There are some industry observers who think that the PC world has caught up with the Mac and that by the end of this year, all the Mac's innovative advantages will have disappeared. For example, David Bunnell has pointed out that "Faster chips, more memory, and better packaging are nice, but the creative advantage Apple enjoyed originally is rapidly dwindling." and that "Apple should have been a moving target; instead, it's been a sitting duck." Declaring that John Sculley's much-publicized Knowledge Navigator is nothing more than a PR gimmick, Bunnell is rather harsh in his critical conclusion: "the company has not shown much imagination when it comes to going further with its platform"; "Apple has had its moment in the sun, but it could soon get cloudy in Cupertino."

Is David engaging in another bout of "Apple-bashing" that was the favorite sport of the likes of Jerry Pournelle five years ago? If so, we can best ignore him, but David has been around personal computing long enough. He knows what he is talking about. His charge that Apple should have enlarged the user base by either introducing the cheaper models of the Mac or by licensing other companies to clone it is not entirely groundless. Nonetheless, the main thing to remember is that the jury verdict is not yet in on the OS/2 and the Presentation Manager. We will have at least one or two more years in which to prove that the Mac is not going to be "a sitting duck." New software in AI that may emerge out of your work can open a new window of opportunity for the Mac in the face of OS/2 and the PM.

I strongly believe that the future of the Mac in artificial intelligence is bright. Apple has so far played its AI card very well. Additionally, the reorganization of APDA and the continuing support for MPW and other related software development have been impressive. In the meantime, you and I, as AI programmers for the Mac, will have to carry our share of the load, which may be heavy at times but will be very exciting. Do you know why I just said the load may be heavy? The documentation for Symbolic's MacIvory fills ten volumes!

One encouraging news is that many astute AI industry observers look favorably upon Apple's entrance into AI industry. For example, Philip Chapnick, the editor-in-chief of *AI Expert* , wrote in his editorial in the November issue, 1988: "AI industry will have to pay attention to Apple from now on. Keep an eye out for Apple; I predict the company will stun AAAI in the next few years."

Mr. Chapnick and other AI commentators are mindful of the AI tools for the Macintosh that have been announced one after another. With the advent of Texas Instruments' microExplorer and Symbolic's MacIvory, a growing number of PC users are attracted to the Macintosh as the leading AI workstation of choice. The Macintosh version of GoldWorks II, developed by Coral Software is already shipping. The fact that Apple acquired Coral LISP shows how serious Apple is in moving into AI. In the meantime, Neuron Data's Nexpert in conjunction with Oracle will be blazing new trails in combining the work of AI and data-base programs, far surpassing the older and well-established rivals.

Throughout 1970s it used to cost several million dollars to buy a mainframe machine big enough to run any AI research program. Dedicated LISP machines started to become available in the early 80s, costing between $100,000 and $150,000, thus making AI research more affordable to more people and institutions. The cost for doing research with LISP has come down a lot lately. Right now the Sun-4 with its SPARC CPU will cost almost $80,000 with the LISP package. Texas Instrument's microExplorer lists for $25,000 and is much cheaper than other LISP machines. Symbolic's MacIvory will also become a strong competitor costing around $30,000. Over all, as Joseph Somsel pointed recently, the Macintosh II, priced at about $6000, rates as a best buy among many personal computers available now for AI. The ParcPlace benchmarks for Smalltalk-80 showed that the Mac II beats some professional UNIX workstations from Apollo and Sun.

Both microExplorer and MacIvory use a NuBus board microcoded to facilitate the dynamic handling of data types, and garbage collection. MacIvory's processor was designed to support Symbolic's Favors subsystem, which is an object-oriented extension to LISP. The microExplorer interfaces to the Macintosh Toolbox by way of object-oriented extensions similar to those of native LISPs such as Allegro. On the other hand, MacIvory uses a C-level interface for linking toolbox calls with the 68020 and provides a full complement of predefined Macintosh Toolbox traps.

The microExplorer software comes in three convenient ways: a runtime system software; a development system with editors, debuggers, and an inspector; an advanced development system that includes the source code in LISP. MacIvory comes in two packages: one for delivery, the other for development, including the source code, a new garbage collector and a new scheduler. In short, the AI researchers using microExplorer and MacIvory will be able to work in the LISP dedicated work environment through Sun microsystem's remote procedure call (RPC) protocol, while using the general-purpose software of the Mac.

When IBM offered an AI-like retrieval program called "Intellect" to be used on top of databases, the number of incoming requests for information increased by the factor of 10 or more. This big jump in the retrieval requests has had a pleasant surprise for the IBM sales division--the first $10 million in "Intellect" sales generated more than $30 million in the hardware sale. In short, artificial intelligence is a great way to sell more computers!

5. LISP and AI Research

LISP's popularity in AI research comes from such features as easy and flexible symbol manipulation, automatic memory management, sophisticated editing and debugging tools, and the uniform treatment of program codes and data, which means that a LISP program can modify its own codes easily as its own data. This last feature lends itself to writing programs that learn new rules or modify existing ones in the knowledge data base. Increasingly, LISP is becoming an all-purpose programming language, not strictly limited to applications in AI. At MIT and Carnegie-Mellon, LISP has been used for years as the primary language in the basic introductory courses in computer science. Some people have even argued that an understanding of LISP is de rigueur for computer science majors.

In a sense LISP is an old language since its history can be traced back to the 1950s, making it nearly as old as Fortran. LISP (from LISt Processing) was created by John McCarthy from a powerful but difficult language called IPL. LISP at the beginning was very slow and required a lot of memory, and no demand arose for any early standardization for commercial purposes.

Freed from immediate exploitation and sheltered in academic circles, LISP innovations flourished and cross-fertilizations among numerous dialects continued over the years. LISP in this fashion has been allowed to mature, giving it vigor and vitality. Many of the new features that have evolved are now added to COMMON LISP, the first standardization that was adopted in 1984 by a consortium of industrial and university LISP programmers. COMMON LISP standards do not cover graphics or machine-specific user interface modes but make newly written LISP programs portable over a wide range of computers. If you are new to AI and lack a background in this lan guage, I would highly recommend *LISP* by Winston and Horn published by Addison-Wesley.

ExperLisp was developed four years ago by the ExperTelligence of Santa Barbara, California and it was the first LISP available for the Mac. Borts and Diamant published some interesting benchmarks for IQLISP and muLISP, two LISP implementations for the IBM PCs. More than two years ago I used three of their benchmarks on ExperLisp, which, not surprisingly, surpassed these two IBM competitors rather handily. The following are the results of my short test:

Benchmark 1:

```
(DEFUN FACT (X)
    (COND ((LESSP X 2) 1)
            (T (TIMES X (FACT (SUB1 X))))))
```

RESULTS:

VERSION	TIME for running FACT (500)
IQLISP	19.0 (seconds)
muLISP	14.5
ExperLisp	7.2

Benchmark 2:

```
(DEFUN FIB (X)
    (COND ((LESSP X 2) 1)
    (T (PLUS (FIB (SUB1 X)) (FIB (DIFFERENCE X 2))))))
```

RESULTS:

VERSION	TIME for running the Fibonnaci series to the seventeenth term recursively.
IQLISP	23.0 (seconds)
muLISP	8.5
ExperLisp	3.6

Benchmark 3:

```
(DEFUN REV*  (X)
        (COND ((NULL X) NIL)
              (T (APP (REV (CDR X)) (LIST (CAR X)))))))

(DEFUN APP* (X Y)
        (CON ((NULL X) Y)
             (T (APP (REV (CDR (REV X))) (CONS (CAR (REV X)) Y))))))

    (APP '(A B C D E F G) '(A B C D E F G))
```

*{'REV' and 'APP' are slower versions of 'REVERSE' and 'APPEND' and provide a good test of the recursion speed of three LISPs.}

Results: VERSIONS TIME

IQLISP	240.0	(seconds)
muLISP	30.6	
ExperLisp	10.3	

It should be noted that PERSONAL CONSULTANT, a knowledge engineering language from Texas Instrument, was written in IQLISP. As you can see, it comes out of our tests the slowest. In the third benchmark program, for example, IQLISP runs about 24 times slower than ExperLisp! And do you remember the Campbell Soul Company's Sterilizer expert system that I talked about earlier? That system was implemented by Texas Instrument with this slow IQLISP. Can you imagine how much faster the system would run in the Macintosh II?

I understand that Mr. Cimino, the retiring expert on the Campbell sterilizer, is currently redesigning the entire cooker environment. Then, the existing expert system will have to be replaced by the time new cooking equipments are set up. Thus, there will be an opportunity for the Mac-based expert system to replace the old expert system. Many other windows of opportunity will open up as more AI researchers start to write the Mac-based expert systems.

6. Ordinary Language and Common Sense

Programming computers to understand ordinary language and common sense is one of the most difficult challenges facing AI researchers today. AI researchers are having a hard time figuring out how to make computers understand metaphors and analogies. In order for a computer to deal with a sentence, it must be translated into an internal representation that captures its meaning. But computers cannot deal with vague or ambiguous words. For example, a computer program cannot cope with sentences such as "He is a pig." or "John is yellow." without knowing a lot more about ordinary language and specific contexts in which the statements are made.

In 1975 Minsky proposed 'frames' or "schemata" as the device to reflect the structures embedded in our words. He began by showing that AI researchers must understand and capture "common-sense thought" more accurately and that most theoretical studies in AI and in psychology have been too fine-grained, local, and unstructured to account for effective common-sense thought. According to Minsky, knowledge is not just a collection of separate simple fragments but is based on frames and "chunks" of reasoning.

This concept of frames was further refined into 'scripts" by Schank and Abelson. For example, to understand the two sentences, "John sat down in the restaurant. The waiter took his order.", we need more detailed information about restaurants, the role of waiters, and what ordering means. Scripts are ordered sets of information about stereotypical cultural events such as eating in a restaurant, visiting a doctor's or a dentist's office. Since scripts turned out to be too large, Schank later reformed them so that they would describe smaller units such as the paying scene, which can be applied to both visits to the doctor's and dentist's offices. This is how he proposed MOPS (Memory Organization Packages) as a better units of stereotypical situations. More recently, a multitude of other theories and explanatory hypotheses such as case grammar, conceptual dependency theory, conceptual graph, and semantic networks have been proposed.

If you find something interesting along this line, you should read the following two books: *Scripts, Plans, Goals, and Understanding* by R. C. Schank and R. P. Abelson (Erlbaum, 1977) ;*A Theory of Learning in Computers* and People by R. Schank

(Cambridge, 1982). Those of you who want a more recent reference book, I would highly recommend *Approaches to Knowledge Representation* edited by G. A. Ringland and D. A. Duce (John Wiley, 1988).

Even if you are just starting to take interest in AI, a good thing is that after even a few weeks of intense reading and studying, you will be able to do some first-hand research in any area of your choice and can make a significant contribution, whether the area be in knowledge representation or any one of fields such as semantic network, fuzzy logic, rule-based systems (including production rules and expert systems), analogies, metaphors, etc. You can choose LISP, PROLOG, or any other languages such as C, C++. Since research projects in vision and robotics will require more expensive hardware, you should be aware of the additional expense.

Apart from such original research projects, you can start doing some interesting AI programming with expert shells that are readily available such as MacSmarts and Instant Expert, as I will demonstrate in the second part of this morning's session. You will be amazed by what you can do with these off-the-shelf development environments. If enough of you go to work in AI, perhaps we can have many AI sessions at the fifth MacHack Conference next year.

If you hear someone say that there has been no significant new idea in AI since the late 1970s, you shouldn't believe a word of it and ask the person making the statement where he or she has been in the last ten years. For AI research has been a highly innovative and productive field during the same period. I can recite easily more than twenty significant new ideas that have emerged in the same period. Ongoing research projects are resulting in so many seminars and special conferences, both domestic and international, that it is hard to keep abreast of even major publications. If you have access to the UseNet, you should take a look at all the lively and at times heated discussions going on everyday among AI researchers not only from this continent but from Europe, East Asia, and Australia.

7. New and Old Research Projects in Artificial Intelligence

In science we usually expect that initial success achieved on a small scale can be extended to larger systems. Yet in AI, promising techniques and methods that seem to work well within tightly constrained domains fail or fall apart when applied to wider contexts or "real-world" situations. The late Y. Bar-Hillel called this the "fallacy of the successful first step." For example, Winograd's natural language understanding program (SHRDLU) created quite a stir in 1972. But this program as well as other "story summarizing" programs developed in Schank's lab at Yale failed to live up to their promises. Computer summarization is turning out to be as difficult as machine translation.

Those who believe that computers are nothing more than a big adding machine argue that computers cannot be made to be intelligent or to be able to learn anything new. Stephen Wolfram is one of those who believe otherwise and has proposed a very interesting theory. He contrasts traditional engineering design with the emerging art of complexity engineering. In a traditional artifact, the engineer proceeds from a detailed logical description of every part of a system and all its behaviors. The system, if successful, does what it is designed to do---nothing more and nothing less. The parts of the system interact with each other in tightly constrained, often linear, ways. Motions are usually periodic and synchronous. Failure in one part of the system often causes catastrophic failure of the entire system. The system can usually tolerate only a limited degree of environmental change.

A complex system, on the other hand, consists of a large collection of individually simple parts, each having only a limited repertoire of possible behaviors. Each part interacts with its neighbors according to a fairly short list of typically nonlinear transition rules. The system as a whole exhibits enormously complex emergent behaviors, which the "designer" usually cannot predict in detail (since the system is computationally irreducible). The complex system can potentially exhibit other desired properties---e.g., robustness and adaptiveness. The trick in complexity engineering, of course, is to select the transition rules that yield the desired emergent behaviors. We are only just beginning to learn how to do this. The current study of neural network, including the Brain Simulator, is expanding rapidly and will shed new light on the rest of AI research.

The human brain, in and of itself, is not intelligent and does not understands anything. A memory understands nothing in itself or by itself. But put many neurons together, set them in motion, and consciousness and intelligence emerge from the interactions of memories, sensory input and output organized in our brain. Similarly, the computer simulation program does not understand anything. Neither does the individual component of the computer. But put all together in the right way, set them in motion, and some rudimentary consciousness may emerge from the interaction of all the parts. It seems to me that Wolfram is one of those who want us to move into this direction of research. Among others, Jerry Fodor has been advocating this line of research. He is approaching AI from the perspective of philosophy , as is true in my own case.

Thus, it may be misleading to say that computer intelligence is limited because a computer can do only what it is programmed to do. If Wolfram and Fodor are right, then computers can be programed to learn and be creative. New and unforeseen capabilities may be generated out of simple components. In this regard, some of you may be interested in doing some research with the new field of chaos study. This will be a new and fascinating area that will have a far reaching scientific and philosophical implications.

Before we move to the next section, I ought to mention one AI topic which you may look into, namely vision. A normal person acquires more information through his or her eyes than in any other way. Children learn vocabulary by seeing and then pointing (dog, house, shoe); they learn physical skills by imitating what they see. Thus Vision is the most important human sense and has been studied ever since the beginning of the computer science. The results have been frustrating and poor. Some AI researchers were at first naive enough to believe that they could solve the problems related to vision in a summer or two. Things have not turned out that way. That is why a cynic commented: "There is more to vision than meets the eye." Be that as it may, *Vision* by D. Marr (Freeman, 1982) is a good comprehensive introduction to this field of AI research.

If you are new to AI, you should start with *Introduction to Artificial Intelligence* by Eugene Charniak and Drew McDermott (Addison-Wesley, 1987). This is much more comprehensive and recent introduction to the field than the earlier, classic *Artificial Intelligence* by P. H. Winston (Addison-Wesley, 1975). If you are new to expert systems, the book to read to start off is *The Handbook of Artificial Intelligence* co-edited by A. Barr and E. Feigenbaum (Kaufmann, 1982). If you are interested in multi-value logic and fuzzy reasoning, I recommend very highly *Expert System and Fuzzy Logic* by C. V. Negoita (Benjamin, 1985). An outstanding introduction to the expert system is *Expert Systems* by P. Harmon and D. King (Wiley, 1985). Specially these last two books have very extensive bibliographies at the end that will guide your further reading and research.

In looking around for an interesting and also profitable research idea, you ought to remember that human creativity is often based on unforeseen combinations of old ideas. Henry Ford combined the idea of an automobile with assembly-line process pioneered by Sears Roebuck. The electric bulb was born when Edison enclosed a carbon thread in a vacuum to prevent it from burning up. Vacuums and carbon threads were studied for a long time before, but it took Edison to combine both for something new.

Spreadsheet programs such as Visi-Calc and Lotus 1, 2, 3 came from an old AI idea known as "constrained propagation." Dan Bricklin and Bob Frankston packaged the idea of constrained propagation behind a matrix of cells that looked like multicolumn accounting sheet. When accounting and financial specialists welcomed Visi-Calc because it speeded up their work, a new industry was born. I need not go on to tell you the story of how Mitch Kapor had some ideas on improving Visi-Calc and came to market Lotus 1, 2, 3 on his own, thus becoming a multi-millionaire. There are more than a few AI ideas simmering around the AI labs that will make some enterprising young hackers rich. If you are interested, I will be glad to talk more about them.

8. The Educational Implications of Artificial Intelligence

AI has major implications for education and work force of any nation. Trends in computer control, personal computing, and networking imply that large organizations, whether commercial, governmental, or military, will increasingly use information systems that are widely distributed. Improving performance through distributed information will therefore depend in large part upon the education and training of workers. So in the long run educational excellence is vital to the future of national security.

For over a decade, the alleged weakness of American education, as reflected in the sagging SAT and other test scores, have been widely voiced and discussed pro and con. American education has been forced into a fundamental self-examination about its educational methodology by the dramatic industrial and technical achievements of Japan. As American people have realized the crucial role played by education in the Japanese success and their own lag, the stir and ferment of the current debates on education in America have indeed spread far beyond the boundaries of professional circles.

What has all this got to do with the Macintosh? Several years ago Anderson and Reiser of Carnegie-Mellon made a detailed study of how our students learn mathematical, logical, and programming skills. They found out that our students, if helped by private human tutors, needed only 11 hours to learn as much as what other students learned in 43 hours in a class-room. Since the major role of the tutor is to make the problem-solving episodes more effective learning experiences, Anderson and Reiser have been trying over the last five years to implement a LISP tutoring program. They have found out that a computer-based tutor is almost as effective as a human tutor.

At the time when both professors released their findings in the *BYTE* (April, l985), their LISP tutoring system required a VAX 725 with 3 mega-bytes of memory for each and every one of students. This was too costly for the tutorial program to be used widely and effectively. Since then, however, the plummeting hardware prices and the arrival of faster computers promise to give us low-cost, reliable tutorial computers that will be more affordable and cost-effective. If the Macintosh computers can be made to work with the LISP and other tutorial programs not only for colleges but also for grade-level science and mathematics, the prospect is real and exciting for providing many more students with computer-based tutorial programs. When this goal is reached, the resulting benefit to American education and society at large will be great indeed.

To give you another example of what expert systems can do for American education, let me tell you what has happened to the St. Louis Public School System. Three technicians and two PC experts used to rush around in the city school system to take care of 2,000 PCs at 160 different locations. But a consultant wrote a hardware diagnosis system and technicians make 30% fewer service calls. There are a number of other expert systems that have been developed for teachers, administrators, and students. One book that I highly recommend here is Artificial Intelligence and Human Learning edited by John Self (Chapman and Hall, 1988). I reviewed this book in the December, 1988 issue of the *BYTE* magazine.

As many of you know, the United States is already defended with the aid of many military "expert" systems, ranging from HASP/SIAP that can detect and identify ocean vessels by interpreting sonar sensor data, to TATR that helps Air Force targeteers develop plans for attacking enemy airfields. It is about time that the resources of artificial intelligence be brought to bear upon the educational needs of the nation as well. Equipped with evolving AI tools, the Macintosh can and should play a vital role in this revolutionary development.

9. The international Implications of Artificial Intelligence

As our lives become more dependent on technology, we are becoming more vulnerable to even a slightest miscarriage of present plans or proposed innovation. What is the point at which our explosive growth will overwhelm the natural checks and balances in the ecological system? How long will the earth sustain such growth, absorb our wastes, while still remaining a viable habitat for humanity? Accustomed to slow and sporadic changes, we are still captive in much of our thinking to the obsolete concepts and categories.

The scale of our global system of material production and distribution, and of communication and transportation has now grown beyond the capacities of a single national or regional group to sustain and operate on its own. These systems are interdependent upon the resources of the entire planet, not only of raw materials and products, but also of experts and specialists. The spirit of independence, self-sufficiency, and the "manifest destiny" are concepts that were useful once when nation-states were isolated.

Traditionally the function of the military was to achieve victory over the enemy through armed conquest. The cost of military defeat was to surrender. Now the cost of even a successful military attack is MAD (Mutually Assured Destruction). That is why Herman Kahn once declared that "If these buttons are ever pushed, they have completely failed in their purpose." In short, the nuclear bombs are useful if and only if they are never used. When he was told of the atomic bomb dropped on Hiroshima, Albert Einstein was reported to have said: "Ach! The world is not ready for it; the unleashed power of the atom has changed everything save our modes of thinking, and we thus drift toward unparalleled catastrophe." If the tools of AI can make us realize the requirements of peace and what we must do collectively to save us from the atomic bomb, then there may still be some slight hope for the future of mankind. Einstein's gloomy prediction may turn out to be just that. As Reinhold Niebuhr used to say, "If our hopes are dupes, our fears may turn out to be liars."

In *The Imitation of Christ*, Thomas A. Kempis observed that "All men desire peace, but very few desire those things which make for peace." Mere absence of war is not a stable condition of peace. True peace requires more than mere con-existence of basically antagonistic nations. Peace calls for extensive cooperation of nations in the pursuit of common values of humanity. Former Secretary of State Kissinger pointed out that the instability of our world stems at "its core from a philosophical schism which makes the issues producing most political debates seem largely tangential," and that, although such philosophical differences in style and perspectives are not unprecedented, "what is novel is the global scale on which they occur and the risks which the failure to overcome them would entail." He continued as follows: "The challenge of our time is whether we can deal with consciously and creatively with what in previous centuries was adjusted through a series of more or less violent and frequently catastrophic upheavals. We must construct an international order before a crisis imposes it as a necessity".

I believe that AI tools and expert systems can play a big role in the emergence of such an international order. Further research must be undertaken to create more expert systems, which may make it possible for the world governments and leaders to deal with a growing list of global challenges, ranging from apartheid and AIDS to the greenhouse effect and acid rain. You may say: fine, so what is the solution? How do we do it?

The expert systems that we will need can be modeled after a specialized expert system called Picon developed by GigaMos Systems not long ago to monitor and control real-time processes. Power plants and chemical refineries are collections of large subassemblies that interact in complicated ways. If something goes wrong in any part of the plant, its variables drift out of tolerance and alarms are sounded for those variables. If the rest of the plant works normally, the problem can be easily fixed. But the trouble usually spreads beyond the point of its origin. At times, the first alarm may come from a part of the plant that is normal but that has been disturbed by a malfunction somewhere else.

I understand that alarm processing can get exciting--a nuclear plant may generate as many as 800 alarms within the first 2 minutes of a major accident. The control panel with 3000 lights may lit up all over at once. Of course without paying attention to the sequence in which red lights come on, there is no way to find out what caused the accident to do anything about it. Now Picon is designed to remember the order in which alarms occur and uses its process knowledge to advise the operator what to do.

Picon cannot scan 20,000 variables that are in the plant fast enough to analyze all the alarms, so it uses a process called "focusing" to decide which data to collect. If the major variables indicate that a part of the plant is not functioning well, Picon reads the rest of the variables for that subsystem and uses more detailed rules to see whether there really is a trouble. Without Picon, operators would have to respond to every variable that drifted out of tolerance and decide whether or not to intervene.

Two years ago almost all the nations in the world signed the international agreement to reduce the consumption of materials destructive of the atmosphere's ozone layer that at the national are profitable but that at the global level are dangerous. Would we not be in need of an expert system like Picon to enforce the treaty on a world-wide scale sometime in the future? If your AI research with the Mac can pave the way for such expert systems, you will deserve the Nobel peace prize!

Our world is becoming interdependent environmentally as well as economically. And this interdependence has lessened the capacities for unilateral action of all na tions. But political structures around the world are still rooted in the centuries-old concept of the nation-state. Perhaps the future leaders of the world can play Version 20 of Chris Crawford's *Balance of Power* running on the Macintosh 20CX in Year 2010 and learn to manage the political affairs of the world a little better, thus preventing the nuclear war. If so, then the Macintosh would have played a great role in saving not only "the rest of us" but the whole of humanity.

In a book entitled *The Art of Discovery*, the famous German philosopher, Gottfried Leibniz, made the following observation: "The only way to rectify our reasonings is to make them as tangible as those of mathematicians, so that we can find our errors at a glance and when there are disputes among persons we can simply say: Let us calculate, without further ado, in order to see who is right." If expert systems can be developed for understanding international conflicts and for removing causes of wars before the outbreak of violent upheavals, then Leibniz's hope for the triumph of logic and reason in human affairs would be vindicated. If the retiring engineer at the Campbell Soup company can serve as the domain expert in order to produce delicious soup with the use of a cooker expert system, I see no inherent reason why someday great diplomats such as the Henry Kissingers or George F. Kennans of the next century cannot be similarly consulted for their experience and insight in the worthy cause of peaceful resolutions of international conflicts. Given the limits of our knowledge of rule-based systems, it will be a while before we can actually use such a system at the United Nations. No matter how difficult it may be to build such a system, it is certainly not a wishful thinking nor a pipedream. Let us remember the UNESCO charter that proclaims that "War starts in the minds of men, and the answers to war must also be found in the minds of men."

10. Three Case Studies of AI Software for the Maintosh:

(1) Go Explorer by Anders Kierulf

(2) Instant Expert Plus by Human Intellect Systems

(3) Cognitron by Cognitive Software