THE 2001 SIREN'S SONG: AI IN THE COMMERCIAL MARKETPLACE

Sharon Bishop
Information Technology Group
Hewlett-Packard
11000 Wolfe Road, Cupertino, Ca., 95014

Summary

This paper gives a brief overview of Artificial Intelligence (AI) and the emerging market parameters that are motivating the use of AI, how AI differs from traditonal programming, examples of tools we take for granted today that are the result of the "fallout" from AI research over the past 20 years, expert systems and HP's activity in AI and finally, some trends and future predictions we might anticipate for this technology.

Introduction

Is Artificial Intelligence the modern-day sea nymph -- part bird and part woman -- of Roman mythology who lured sailors to their death on rocky cliffs by seductive singing? Are we the modern-day sailors snared by the eager contemplation of songs foretelling the day we will have machines that eliminate the need for highly paid human experts, computers that speak and understand speech, robots and vision systems that have the ability to recognize patterns of input stimuli and act accordingly, learning systems that learn from usage patterns and modify themselves consequently, self-programming systems, etc.?

Artificial intelligence, popularly defined as the science of enabling computers to reason, make judgments and even learn, has generated considerable interest in the computer industry over the past year. Optimistic analysts forecast a market growth in excess of $5 billion by 1990. Stories about fifth generation computers and the Japanese effort to build "thinking" computers by 1992 have added urgency to competitive efforts in the U. S. and Europe.

The evolution of AI can be viewed as a two-part process. The first stage of this evolution is the 20-year old ongoing research effort aimed at developing computer systems that can emulate human thought and decision- making processes. The second stage is the newly emerging commercial market for AI-based products.

Hewlett-Packard has, for some time, been developing AI technology at our Corporate Research and Development laboratories. Our first entry into the AI market was presented at the International Joint Conference on Artificial Intelligence, held in Los Angeles during August 1985. Initially, we are addressing the productivity issue, that is, the motivation for power tools for programmers. Other examples of internal development activity using our AI technology include expert systems for natural language understanding, an organic compound analyzer, an I/C photolithography advisor, computer peripheral diagnosticians, and an electrocardiograph advisor. HP views AI not as a revolution, but as part of an evolution toward more powerful and sophisticated computer systems, and we have targeted AI technology to play an important role in our future.

Definition of Artificial Intelligence

Artificial intelligence is a buzzword for a group of technologies It is a multidisciplinary study, utilizing and synthesizing knowledge from several fields including cognitive science, philosophy, sociology, computer science and electrical engineering. It is not in itself a commercial field but a science and a technology. As an academic discipline, it is a collection of concepts and ideas which are appropriate for research, but which cannot be marketed. However, artificial intelligence is the scientific foundation for several growing commercial technologies. AI techniques make a variety of AI applications possible, including expert systems, computer-aided instruction, automated programming, natural language understanding, vision systems, voice recognition systems and robotics. The unifying factor in all of these diverse applications is the encapsulation of some knowledge or expertise related to humans.

Goals of Artificial Intelligence

The primary goal of Artificial Intelligence is to make machines smarter. The goal of the basic researcher is to understand what intelligence is, and the goal of the entrepreneur is to make machines more useful. To date, we have used computers to evaluate problems we have already solved, for example, computers give us solutions to models and simulations or they solve algorithms for us. The goal with AI is to move toward having computers participate in the analysis and synthesis of problem-solving -- to help us develop the solutions, rather than simply calculating the results. In it's short history, AI as an academic discipline has developed techniques that support computer systems that do seem to exhibit intelligence. For example, researchers have developed systems that comprehend natural language, systems that can outperform human experts, systems that see and robots that manipulate tools. The caveat of course, is that the current systems are limited and work successfully only in a narrow and specified domain. They cannot respond to unanticipated events in the environment and learn efficiently from experience. Still, even though there is criticism of the current systems and their limitations, they have demonstrated significant capabilities that have captured the imagination and attention of commercial and government institutions. Rapid developments in more powerful computer systems at lower costs have increased the number of applications that are now feasible and economical. Today's challenge for commercial users is to identify practical, worthwhile applications and to develop and market these applications.

AI Systems vs. Traditional Systems

Artificial intelligence development is usually characterized by the use of tools and techniques which facilitate the handling of variable, ill-defined and unspecified programming parameters. While the presence or absence of such tools and techniques does not necessarily make a product "genuine AI", it is clear that AI places new demands on traditional tools and languages and the programmers who use them.

Generally the processing power and memory demands on AI development has led to a dedicated workstation approach. The terms "symbolic processing", "object-oriented programming," "knowledge-representation" and others refer to the process of representing information in a fluid rather than highly structured form. These techniques are often required because in contrast to traditional programming objectives, data must be represented symbolically and related through parameters and conditions in a program similar to the human decision-making process. Languages like LISP, PROLOG and SMALLTALK which are highly interpretive and are function or object-oriented, as opposed to procedurally oriented, lend themselves to incremental programming and rapid prototyping. In contrast, conventional programming techniques are ill-suited to handling uncertain or changing specifications. Virtually all modern programming methodology, such as structured design is targeted to ensure that the implementation follows a fixed specification in a controlled fashion, rather than wandering off in an unpredictable direction. In a well-executed conventional implementation project, a great deal of internal rigidity is built into the system, ensuring its orderly development.

AI applications in general incorporate a knowledge base and separate units of code that operate on the knowledge base, the combination of the two being used to solve a problem. Knowledge used by the system is represented explicitly within the system. In constrast, traditional computer programs state procedures explicity, and knowledge about the domain is implicit in the procedures. This concept of the difference between traditional computer systems and AI systems forms a useful basis to distinguish AI programs from traditional programs. The following table highlights some of the differences between traditional systems and AI systems:

Traditional Computer Systems vs. AI Computer Systems

TRADITIONAL SYSTEMS	AI SYSTEMS
Host/timesharing approach	Workstation/PC approach
Separate programming tools, usually not well integrated	Integrated environments, sophisticated programming tools
Minimal prototyping; systems are hard to modify as the environment changes	Rapid prototyping; systems are easy to update, treat knowledge as data
Tend to degrade with age	Can improve with experience increases; knowledge explicit in system and easily updated
Use compiled-only languages, follow complete predefined step-by-step procedures, algorithms	Use function or object-oriented languages, fluid design, follow a line of reasoning
Process numbers	Process symbols, as well as numbers
Use a "black box" approach	Can explain line of reasoning

The AI Productivity 'Fallout'

It is interesting to note that virtually all the environments we use today in traditional computing environments, whether tightly coupled to proprietary operating systems, or highly portable accross vendors, take for granted many capabilities that are essentially the "fallout" from AI developoment that took place in the AI research laboratories through the 1960's and 1970's. For example, in the 1960's, the first computer time-sharing systems were developed in AI laboratories to addresss a faster means of debugging very complex programs. Word processing was developed in the AI lab in the 1960's and expert systems started in 1965. Bit map displays, the now-popular mouse controls and the whole structure of display windows were also originated during the 1960's as well as object-based programming (as in SMALLTALK) and a number of other productivity tools.

Commercial use of these tools originally was discouraged by the fact that they require large computing and storage capabilities, but in the 1970's improved computer performance at lower cost made it practical to start using these tools in conventional programming environments. For example, there has been a dramatic improvement in recent workstation and interface technology, with great cost/performance benefits for group work. Powerful networked workstations/personal computers with high-resolution graphics are becoming widespread. Today machines offered by vendors such as Hewlett-Packard, IBM, DEC, Xerox, Sun, Texas Instruments and Symbolics support tools and languages such as LISP and PROLOG at modest cost. These machines provide adequate hosts for radically improved programming environments.

Market Characterization of AI

When looking at the market for Artificial Intelligence, it is important to remember that this market is very young and did not exist in the minds of commercial users even five years ago. There was no market for expert systems or symbolic computing outside the confines of the academic research and development environs. For example, Symbolics, Inc., the leading supplier of LISP machines today, is generally credited for creating a market that nobody knew they wanted; potential users didn't even know what a LISP machine was. Symbolics created a burgeoning industry out of something that researchers did not realize existed, but found to be a very useful tool once they were educated as to its capabilities.

Another point to keep in mind when anticipating the AI market is that you have to view it in terms of commercially viable products that are for sale and can be used by the end-user. There is continuing basic research going on at universities attempting to understand the nature of intelligence and how to replicate the human brain that isn't transferable today into commercial products.

Market segmentation and sizing are two of the hottest topics in AI today. For puposes of this paper, the market will be segmented into the following areas: voice and vision recognition systems, expert systems, natural language understanding and AI computers and languages. Each of these segments is a viable market in and of itself.

THE ARTIFICIAL INTELLIGENCE MARKET
PRODUCT AREAS - 1985

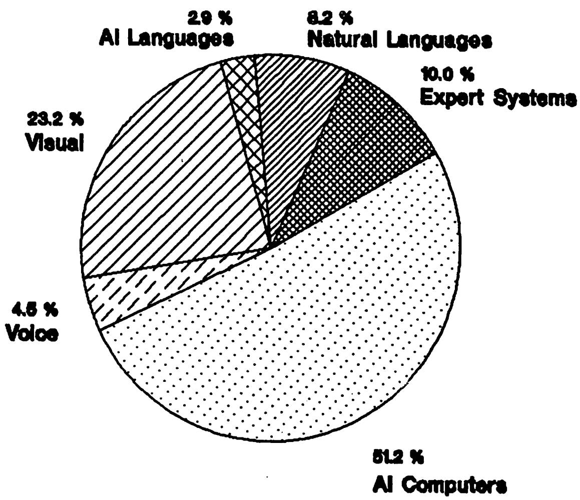

SOURCE: DM Data, Inc.

* AI COMPUTERS

This is the largest segment of the industry because it includes machines that are LISP-oriented (such as those produced by Symbolics, Inc.) and general purpose computers that can also be used for AI applications (such as Hewlett-Packard, Digital Equipment and IBM).

* AI LANGUAGES

AI languages is really a subset of AI computers. This segment includes LISP and PROLOG programming languages and environments for AI applications. LISP is the most popular programming language in the United States. One of the principal features of LISP is that it allows the user to compute with symbolic expressions rather than numbers; that is, bit patterns in a computer's memory and registers can stand for arbitrary symbols, not just those of arithmetic. LISP itself is the programming language often used to develop expert system building tools. PROLOG is the language of choice for AI applications in Europe and Japan. It is based on the idea of representing a program as a set of declarative statements in a form of logic.

* EXPERT SYSTEMS

Expert systems are currently the most visible of the AI technologies because they are the most easily understood in terms of what they actually do: an expert's knowledge in a particular domain is programmed into an expert system and that system emulates his/her behavior. The expert system is composed of two parts: (1) a knowledge base which is the actual information given by the expert and programmed into the system (for example, a doctor who specializes in bone grafts); and (2) an inference engine which is the actual computer program that runs that information, making the correlations and associations between the information that has been programmed in. Expert systems is forecasted to be the largest growing segment of the AI market. More people are aware of its potential capabilities and will plan to utiltize expert systems in order to achieve competitive advantages and profit.

* NATURAL LANGUAGE

Natural language does not encompass voice recognition or actual speaking into a microphone. It is simply communication with the computer or terminal in written conversational style. Natural language systems have to deal with many ambiquities, such as regional slang, and the correct interpretation of a number of variables in sentences. Therefore, it is forecasted that this segment will not grow as rapidly as expert systems. It is expected that natural language will be included in with other packages, such as expert systems and AI hardware. Since natural language still requires the end-user to type, some market analysts forecast the value of this type of interface will be much less than voice recognition.

* VOICE RECOGNITION

Voice recognition tehnology in the context of AI is not to be confused with voice store and forward, which is the recording of a human voice and then replaying it later; or synthesized speech, which is a computer emulating the sound of a voice. It is actually the ability of a computer to respond to vocal commands exactly like it responds to keyboard commands. Utilizing a microphone instead of a keyboard, users would talk to the computer and it would react accordingly. Again, as with natural language, the system has to deal with many ambiguities: speech patterns, pitch, tone and the variability of each word. Most voice-recognition systems today have 1,000 to 2,000 words of memory and are speaker-dependent. There are companies promising to produce voice-recognition typewriters in the near future with memory lists of 10,000 to 15,000 words, however, it is felt that this technology is at least five years from being a truly viable market.

* ARTIFICIAL VISION

In the U.S., artificial vision is mainly associated with the auto industry and robotics. General Motors has holdings in six of the major artificial vision companies. Ford has holding in two others, and Chrysler has holdings it has not disclosed. Quality inspection and areas where it's too difficult for humans to work or the failure differentials are too small for a human to react to or to be able to perceive are a few of the possibilities for this technology. The major research push today is for 3-D vision systems that have acceptable real-time capabilities for commercial exploitation. The market segment is expected to expand and grow rapidly.

THE ARTIFICIAL INTELLIGENCE MARKET
(Millions of Dollars)

MARKET AREA	1985	1986	1987	1988	1989	1990
Expert Systems	74	145	245	385	570	810
Natural Language	59	125	210	320	465	650
Visual Recognition	168	260	370	500	660	840
Voice Recognition	33	55	85	140	200	270
AI Languages	21	35	45	65	80	105
AI Computers	364	510	710	970	1250	1570
Government Contracts*	95	150	150	155	175	200
TOTAL	719	1130	1665	2380	3225	4245

* Not in total; already included in other areas.

SOURCE: DM Data, Inc.

User Needs Characterization

In-house corporate and university research development projects are the largest users of artificial intelligence today. The U.S. government, primarily the military, is the next-largest user.

It is important that we look at the underlying need that is driving the purchase of computing systems that support symbolic computing, that is, that area of computer technology where the aim is to optimize the processing of symbols, representing objects or concepts rather than numerical data or text.

* There is an overwhelming need for a major breakthrough in programmer productivity. Symbolics, Inc. has stated that their customers, with a Symbolics system, claim productivity increases in software development by a factor of between two and 50. Even taking the low end of that range, if a customer can solve a problem in one year instead of two, they can gain a major competitive advantage.

* Symbolic computing allows a customer to pursue applications that are difficult or impossible to approach using conventional computing techniques. These complex applications require a sophisticated development environment. Customers purchase systems because of speed, but equally important is a high-performance LISP compiler and a superior programming environment.

* Symbolic computing can fill the need for automation of nonclerical tasks. Just as word-processing has had a major impact on productivity for clerical personnel, symbolic computing can have a profound effect

on the productivity of professionals doing such tasks as integrated-circuit design and computer systems maintenance.

The majority of users evaluating AI technology want to develop expert systems. They want tools that are easy to use and lower cost systems with competitive performance. They need access to AI technology through training by vendors or consultants. Systems need to be easy to use with superior user interfaces. Standard networking facilities are needed so they can coexist with multiple vendors. Customers want integration with existing languages such as COBOL, FORTRAN and Pascal to allow use of existing code to extend the usability and simplify the maintenance of their conventional programs.

Expert Systems

Expert systems is the fastest growing segment of the AI market today. An expert system is a program that can reason, make recommendations and explain its reasoning, based on extensive knowledge of a given field such as geology or medicine. It can be thought of as a form of "intellectual cloning". An example of an expert system that is well-known is General Electric's CATS-1. This expert system diagnoses problems in diesel locomotive engines. The knowledge and rules for the system were culled from a 40-year vetern at GE who was an expert in locomotive trouble-shooting. In the commercial market, the term "expert system" is used to encompass most knowledge-based systems. All knowledge-based systems today, whether or not they are truly expert, address a narrow problem domain. It is important to remember that today's expert system can only capture expertise that humans already have. We are still far from the goal of building systems that can reason as people do and learn from experience.

Anatomy of an Expert System

The major components of an expert system are the knowledge base, the inference engine, and the user interface. Other components may be optional such as external sources of information residing in external databases and external sensors. The knowledge base consists of facts and rules about a specific domain. The inference engine communicates with the knowledge base to determine what rules may apply to the current situation and it also keeps track of the status of the problem. The inference engine may also have optional access to external sources of information. In some cases when the data in the knowledge base is not definite, a "certainty factor" may be assigned. When certainty factors are used, the inference engine includes procedures for calculating and reporting the level of certainty for answers to problems. Some expert systems may support an interactive dialog with the user who may ask why a certain conclusion was reached or why a particular question is being asked. Some user interfaces for expert systems may include a natural language processor and a graphics editor.

Hewlett-Packard has an internal tool that is a flexible expert systems construction environment that provides advanced reasoning and representation tools for development of applications. It has served as the basis for the expert systems development activity at Hewlett-Packard.

Application Areas for Expert Systems

Generic categories of applications suitable for expert systems development include interpreting and monitoring sensory input, diagnosing equipment, prescribing and implementing remedies, planning and design and prediction systems. Possible application areas cover a multitude of industries such as finance, law, medicine, military, education, manufacturing and service, agriculture and mineral exploration. Many of the actual expert systems that have been built and are well-known were built at Stanford for medical applications.

While many application areas may be candidates for expert systems development, this should only be chosen over traditional programming techniques when the domain knowledge is not firm. When the knowledge and problem is firm, fixed and formalized, algorithmic programs are more appropriate than heuristic ones and should be used. If the knowledge is subjective, ill-defined and partly judgmental, expert systems and hueristic programming are appropriate.

DIAGNOSTIC AND MAINTENANCE: In today's environment, a list of expert systems would most surely include a multitude of diagnostic systems for equipment. Hewlett-Packard has expert sytems in use today in their response centers to diagnose and trouble-shoot data comm lines and disc drives. Other examples include DEC's SPEAR system for diagnosing problems in their tape drives and PRIME's DOC expert system for assisting with troubleshooting on their 750 minicomputer. IBM has DART for disk fault diagnosis and AT&T Bell Laboratories has ACE to manage telephone cable maintenance.

U.S. MILITARY: The U. S. military has targeted a considerable amount of funding for expert sytems projects and has declared AI-based applications critical to upcoming generations of military hardware and software. Examples of expert systems that have potential include a pilot's associate, battlefield advisor and roving vehicles for use in hazardous conditions.

FINANCE: This application area holds great promise for the use of expert systems. The management of money is a real, measurable goal and those companies that do it best are going to have a competitive advantage. Areas proprosed for use of expert systems are cash management, portfolio management, underwriting, loan management and corporate finance. Financial expert systems would be able to handle huge amounts of rapidly changing data and apply the rules of finance to assist in decision making. For example, when to go public, covering of foreign exchange exposure and so forth. Cognitive Systems, Inc., TAD expert system assists in tax preparation.

MEDICAL: The medical field has also been a natural application for AI because much of a doctor's skill involves applying rules to evidence presented by a patient. HP is marketing an ECG electrocardiograph expert system that aids physicians in the diagnosis of heart disease. The product supports diagnosis justification, modifiable rules, stand-alone use and integration with other systems. MYCIN and ONCOSYN and PUFF are expert systems from Stanford University. MYCIN diagnoses meningitis and other blood-borne bacterial infections and recommends theraphy. ONCOSYN recommends drug treatment for cancer patients. PUFF is used for pulmonary function test interpretation.

MANUFACTURING: This is fertile territory for expert systems to be the brain for robots and to assist managers and engineers in decision-making with massive amounts of data. Hewlett-Packard has a Manufacturing Research Group in our applied research laboratories at Corporate headquarters that is specifically addressing advanced technologies for the manufacturing arena. We are very interested in robotics and machine vision.

SCIENCE & ELECTRONICS: Activity in many of the major electronics firms such as Hewlett-Packard, IBM, Xerox, Texas Instruments, DEC and others bear out predictions that increased importance is being placed on AI technology for in-house use. HP has developed an IC photolithography advisor that diagnosis fabrication defects and recommends correction action. The system allows for single or multiple faults, and gives conclusions and justification and can display video images to help identify faults. HP also has an IC design tool in the form of a procedural langage embedded in LISP for designing integrated circuits from library parts. VLSI is a ripe area for expert systems since we are rapidly approaching the density of one million transistors on a single chip. To design this manually is like designing a city block by block - the complexity goes up exponentially and the opportunities for error are great. Using AI techniques allows the design engineer to take a much more modular approach to design - much like writing a program with the added flexibility for change and prototyping providing definite productivity gains. Other examples of expert systems in use by electronics firms include DEC with their R1 (alias for XCON) expert system to help configure computers and their XSITE expert system to recommend computer site preparation. Xerox has KBVLSI which aids in VLSI design, Daisy Systems markets GATEMASTER and LOGICIAN for performing gate array layouts.

OTHER: Of course there are many other areas that have potential for use of expert systems. In HP, we have internal research applications underway for a number of internal tools, including natural language understanding, expert systems for the office and expert systems to provide more powerful and sophisticated programming environments.

Future Trends

The AI market milestones from now through the 1990's seem relatively clear. The industry and university collaboration that has already lead to development of several expert systems and the emergence of hardware tools and development environments will continue. More sophisticated software and hardware tools will be introduced in the near term and hardware prices will continue to edge down. At the moment, no AI tool supplier can claim to be the leader, but as further refinement takes place in the market and tools, companies will start to turn from "research- driven" to "market-driven" to get their AI products into the marketplace. By 1987 and 1988 it is felt that "off the shelf" expert systems will start emerging for the consumer. AI will become a normal part of the software market. By 1990 AI capabilities will be incorporated into most products.

However, while AI research continues at a brisk pace, full commercialization faces several potential hurdles, among which are (1) limited industry expertise in building knowledge-based systems, (2) the need for development tools and standards that are easy to use and economically feasible and (3) the need for characterizing the still ill-defined market and addressing critical market issues such as niche definition, pricing, channels of distribution and customer training and support. The period of market definition is likely to continue for the next two or three years. Software metrics also need to be gathered to give a more solid indication of productivity gains, since this technology is likely to involve a lengthy learning curve.

Summary

The world of the mainstream computer supplier is changing and the arbiter of that change is symbolic computing. Symbolic computing is the area of computer technology where the aim is to optimize the processing of symbols, representing objects or concepts rather than numerical data or text. In the past, changes in the 8 and 16-bit microprocessors had profound effects in the market, creating new industries and new competitors. Now the emergence of symbolic computing will have even more impact; not only will new markets be created, but existing markets will be severely impacted over the next few years.

In 1985 some major players in AI -- including IBM, Xerox, Hewlett-Packard, and Texas Instruments -- introduced new AI-related products and reaffirmed their committment to AI research and development. The products included AI programming "shells" to build expert systems, programming languages and tools and high-performance, lower-cost workstations that support AI products. It is anticipated that in 1986, these products will become even more widely available, marking a milestone in the evolution of AI from the research and development laboratories to the commercial marketplace. As the market becomes more clearly focused, all major computer companies, mainframe, minicomputer and microcomputer, will enter into direct competition and will vie to offer systems of comparable functionality. This will elevate the competition to a new level.

For AI suppliers, success in the AI market will not necessarily be determined by previous success or failure in the 16 or 32-bit world. More likley, success will depend on the manufacturer's ability to produce on schedule these order of magnitude more complex systems; the ability to adequately supplement the systems with support for AI programming languages (LISP and PROLOG); high performance mass-storage devices and high-resolution graphics displays; the implementation of a large virtual address space and large amounts of physical memory; and the garnering of third-party software support for languages, development tools and applications.

For the major computer manufacturers such as HP, IBM and DEC, it has become a question of not whether they should develop an AI-based product line, but rather which system processor they should choose and when. As with any technological upheaval, the emergence of symbolic computing on high-performance workstations and PC's give new players the opportunity to get the jump on the existing competitors. However, difficult decisions must be made on the basis of technical, marketing and support issues. Choosing the right market niche and third party suppliers will be tantamount to choosing the right horse. Also, choosing a 32-bit processor family is not enough, as many system manufacturers will choose the same processor family. The issues of system design, price/perforamnce ratios, marketing strategies and software support will make the critical difference between success and failure.

For the potential customers of Artificial Intelligence, it is well to heed the siren's song -- with a clear eye and an alert mind. End users owe it to themselves to become educated about this technology and informed of the activity that is taking place. Many vendors are becoming very agressive in marketing tools and an escalation of this activity will surely take place over the next couple of years. It is important to examine your business for potential applications, keeping in mind that commercially viable products on the market today are far from fulfilling the promise of AI. Researchers are still trying to understand and capture the nature of human intelligence -- we do not have machines today that can reason as humans do and learn from experience.

Biography

Sharon Bishop is a product manager in Hewlett-Packard's Information Technology Group in Cupertino, California, with responsiblity for several products, among which are AI lanaguges and programming environments. During her eight years with Hewlett-Packard, she has held senior programming positions, managed a group of application software engineers and, most recently, moved into product marketing for computer languages.