Critical Technology Assessment of the U.S. Artificial Intelligence Sector

Prepared by

U.S. Department of Commerce
Bureau of Export Administration
Office of Strategic Industries and Economic Security
Strategic Analysis Division

August 1994

For further information about this report, please contact
John Tucker, Senior Trade and Industry Analyst, 202-482-3755
Brad Botwin, Division Director
at
Phone: 202-482-4060   Fax: 202-482-5650
e-mail: bbotwin@bxa.doc.gov

For more information about Bureau of Export Administration Publications, visit our World Wide Web Site at http://www.doc-bxa.bmpcoe.org

CRITICAL TECHNOLOGY ASSESSMENT OF THE U.S. ARTIFICIAL INTELLIGENCE SECTOR

Table of Contents

page

Table of Contents

Table of Contents

EXECUTIVE SUMMARY

Overview

Artificial intelligence (AI) is an emerging technology of strategic importance to the military and of increasing importance to the international competitiveness of U.S. corporations. AI's greatest value is automating and increasing the utilization of expert or organizational knowledge (i.e., knowing what to do with given information or circumstances). Although AI technology is still in its early stages of development, many spectacular success stories in both military and major corporation applications serve as testimonials to AI's potential. The best AI systems save companies and the government millions of dollars a year.

The United States leads the world in nearly all aspects of AI technology, largely due to over 30 years of patronage by the Department of Defense (DoD). Currently, the United States alone accounts for over 60 percent of an estimated $900 million global AI market. Cutting edge (usually very expensive) AI applications, accounting for almost two-thirds of the U.S. market, are normally developed internally by DoD and a handful of major corporations (mostly) in the information technology field (AT&T, DEC, IBM, Apple, Intel, etc.).

More mature and proven AI technologies are packaged and marketed by merchant vendors. Most merchant AI systems are small in scope, and range in price from about $100 to $250,000. An estimated 70-80 percent of the Fortune 500 companies use AI technology to varying degrees. The merchant market, comprising about one-third of the global market, is intensely competitive and innovative, but not profitable. Between 1988-1993, merchant revenues more than doubled. However, vendors came and went by the dozens. About five or six leaders have emerged, but in 1993 only one had sales of over $40 million.

The United States leadership position in AI is eroding as the governments and companies in Japan, and as well as in Western European, working together, have gained ground. In select areas of AI, Japan and Western Europe now surpass the U.S. Two major conditions threaten the country's leading position: 1) the slow rate of AI commercialization relative to extensive research and development expenditures, and 2) declines in Defense research and development funding. Cuts in the military budget have reduced overall AI R&D, shifted research away from basic science and long-term projects, and threaten to delay AI's potential from being realized. The commercial sector will not come close to replacing these lost R&D funds, particularly at the basic research level.

The establishment of an AI experts group, with industry, academic, and government participation (representing commercial and defense interests), is needed to provide focus and direction to AI R&D and commercialization issues.

Background

This AI assessment was undertaken under the Defense Authorization Act for Fiscal Years 1991 (Public Law No. 101-310, Section 825) and the National Defense Authorization Act for Fiscal Year 1993 (Public Law No. 102-484, Section 4215). This legislation requires the Departments of Defense and Commerce (acting through the Under Secretary for Export Administration) to submit reports to the Armed Services Committees of the Senate and the House of Representatives on the status of technologies deemed essential to the performance of current and next generation weapon systems, and crucial to the commercial sector's ability to compete in the global economy.

The goal of this assessment is to provide industry executives and government policy makers with comprehensive information and analysis about research into and production and application of AI technologies in the military and the U.S. economy. This includes analysis of the economic performance and international competitiveness of private sector firms and academic institutions involved in the creation, distribution, and use of AI, and the impact declining defense budgets have on the technology. In achieving this goal, the Department of Commerce's Office of Industrial Resource Administration (OIRA) collected information from the public and private sectors with a survey questionnaire and sought expert advice as necessary.

What is AI?

AI refers to highly engineered computer software programs used to make computers do things that appear intelligent - such as reason, learn, create, understand human speech, or solve problems. As a science, AI studies the nature of intelligence, and tries to make computers simulate intelligent behavior. As a technology, AI is used to automate (or extract and synthesize) knowledge from information and databases. AI has evolved into many specialized areas and approaches. Major areas and sub-areas of AI research include:

automatic programming decision making expert data base systems expert systems fuzzy logic game playing general problem solving intelligent computer-aided instruction	knowledge representation knowledge acquisition logic programming machine learning natural language processing neural networks pattern matching pattern recognition	planning robotics search speech recognition theorem proving uncertainty understanding systems vision

(See Appendix A for definitions.)

AI systems can today: 1) help organizations manage knowledge assets and deal with complexity; 2) help experts solve difficult analysis problems and design new devices, 3) learn from examples; and 4) provide answers to English questions using both structured data and free text.

Major AI Technology Applications

AI is a dual-use technology. The same AI system shells or tools can be readily applied across most military and civilian applications. In terms of sales, the most successful and dominant AI tools to date have included knowledge-based systems, neural networks, fuzzy logic, and natural language systems.

In 1993, the global market for AI systems was estimated at about $900 million. The North American (NA) portion of the global market was roughly $600 million of this total. Slightly over $200 million of the NA portion, or about one-third, was sales and license fees collected by AI vendors (the merchant market). An estimated 20-30 percent of this merchant market was exported. Major AI tools include the following:

Knowledge-based (or expert) systems (KBS) have experienced the most commercial success to date. They are used for diagnostic, scheduling, planning, data synthesis, and tutoring purposes, and for automating manuals and running factories. The NA market for KBS totaled an estimated $350 million in 1993. This included estimates of $171 million developed internally by corporations, about $30 million for off-the-shelf systems and embedded components, and $150 million in AI vendor sales and licensing fees.

Neural networks were long neglected, but are projected to grow at a fast rate in the next decade, and perhaps overtake KBS systems. Neural networks efficiently handle huge databases. They are exceptionally good at pattern recognition. Neural network systems are used by the brokerage houses to predict stock fluctuations, by banks to detect fraud, by insurance companies to appraise applicants, and are finding their way into the factory, where, for example, they can sort fruit at incredible speeds. Neural networks and fuzzy logic systems lumped together totaled $150 million in 1993. Neural net vendor revenues totaled about $26 million, and fuzzy logic vendors' about $6 million. As the numbers attest, the bulk of these systems were developed internally by government and corporations, or used as embedded components in other software programs.

Fuzzy logic is the latest rage, although total sales are still small. Invented in America 30 years ago, the Japanese embraced the technique and now lead the world in nearly all aspects of fuzzy logic. Fuzzy logic is used mostly as a control mechanism in camcorders, anti-lock braking systems, elevators, transmission controls, washing machines, and many other products.

Natural language systems (NLS), particularly speech recognition, have benefitted from significant recent advances in the technology. NLS systems are used to interface between humans and machines, and to allow human-machine communication in English in place of a rigid set of commands. This includes two-way communication. Using NLS, a manager can dictate a letter into a machine and have it printed or simply transmitted (paperless option) across the country. The NLS market was estimated at $64 million in 1993.

AI: A Part of the Emerging Knowledge Automation Industry

AI is not a "stand-alone" technology, and to an extent has lost identity as a distinct product in the marketplace. The technology is increasingly embedded or integrated in multi-purpose software packages where it can, for example, increase the productivity, performance, and user friendliness of the package. As an emerging technology with customized applications across the economy, AI is difficult to define as a distinct industry. Based on what it does, however, AI is combining with several other software technologies and emerging as the *Knowledge Automation Industry*. This industry is an important component of the rapidly expanding information age, playing a unique role in automating the process of converting information into knowledge.

Components of the Knowledge Automation Industry (A software industry that facilitates and improves the ability to leverage the value of an organization's in-house knowledge and experience.)	
Artificial Intelligence	makes it easier to manage knowledge and experience assets
Object Oriented Programming	makes it easier to manipulate code
Computer Assisted Software Engineering	makes it easier to manage software development
Client-Server Networks	makes it easier to connect people and computers in an organization and promotes teamwork

The AI sector can be divided into three overlapping components. These are research, commercialization, and applications. The research component is predominantly funded by the Federal Government. The commercialization component transfers the technology to end-users for a price, or more often, a cost if done internally. The applications component is made up of customers -- the businesses and government agencies that purchase AI technology packaged in a variety of ways. Its true value lies in the increased competitiveness conferred on businesses using the technology. This value is difficult to quantify, but substantial.

Research Component - Total research has hovered just above $200 million per year since 1990. The top five research organizations represent almost 40 percent of all AI research and over 50 percent of total basic research in AI in the United States. In terms of research capabilities, the United States is preeminent, but several foreign schools (e.g., University of Edinburgh) and companies (e.g., Siemens) are the equal of the best research organizations in the United States.

The Federal Government sponsored 75 percent (about $150 million per annum) of the AI research in the United States between 1989-1994. Over 80 percent ($125 million per annum) of the Federal total came from the Defense Department. The Advanced Research Projects Agency (ARPA) has been the champion of AI research and development almost from the Agency's inception in 1958.

In the past few years, pressures mounted within the Defense community to show results amid declining defense budgets. As a result, basic research, which accounted for about half of total research expenditures in 1993, is on the decline relative to applied research and development. Basic research is categorized as high-risk, long-term and is often the first area to be cut when budgets and priorities tighten. U.S. businesses sponsored about 15-20 percent of the research, while funding from foreign sources accounted for most of the remainder.

The share of overall AI research undertaken by universities and firms is about equal. Private internal research, which is more focused on short-term applied research and development, could not be quantified, but may range from $50-100 million. This research is undertaken by major corporations involved in the information technologies (AT&T, IBM, Intel, etc.).

Commercialization Component - The merchant market is intensely competitive in marketing primarily proven AI technologies. Most AI vendors are small and entry into the business is low cost. Because of easy entry, the market is perhaps overpopulated. Most vendors lodge themselves in a "niche" where they try to survive. Many have gone out of business while others continue entering the business. The largest vendor had sales of less than $50 million in 1993.

The overall AI market (merchant and internal development) has a high and low end. The high end or leading edge of the industry develops highly sophisticated AI systems that often merge the best of several AI technologies into a single system. These systems are super-engineered expert systems that in a very real sense exhibit capabilities near the human expert level. Leading-edge AI systems are typically developed for and often by major corporations, or the Department of Defense and other large organizations where the savings and productivity opportunities are the greatest. The development of leading-edge systems frequently includes university or think tank participation.

The lower end of the market, the merchant market, is comprised mostly of AI vendors marketing prepackaged proven AI technologies. The AI tools range from simple to complex, are usually inexpensive and, as with advanced systems, can include several integrated AI techniques. An AI vendor may simply sell or license an empty "shell" (or programming format). Vendors typically also offer development services at the customer's option.

With some notable individual company exceptions, the merchant market has not been profitable. Survey respondents as a group for the years 1989-92 showed profits in only one year, 1990. Total revenue for the four year period totalled more than $224 million, with losses of about $11 million. Much of the loss was due to the collapse of the mainframe business. To supplement the survey data, a review of the financial statements of four major publicly held AI vendors (not part of the survey sample) from 1988-1993 provided a similar picture of the financial difficulties many AI vendors have experienced. Total revenues by the four during the 6-year period were $473 million. Losses on those revenues were $106 million (i.e., a net loss of 22 percent). The financial health of the entire industry should improve over the next few years as market conditions improve. AI techniques are being integrated with other software (client-server, object oriented programming) as these vendors adapt to the changing needs of corporate computing.

Applications Component - Military - The Department of Defense is by far the single largest user of AI in the world. In using AI as a strategic asset and management tool, the military has demonstrated the feasibility of the technology, pushed its development, and, through its use, improved military effectiveness. The military uses AI systems for diagnostics, testing, robotics, target recognition, tutoring, war planning, logistics, nuclear test-ban monitoring, database management, and defense-related manufacturing.

Expert diagnostic systems are used extensively by the military to troubleshoot, and maintain military equipment operational. An important value of these systems is facilitating repairs by non-experts on equipment that otherwise could be out of commission for extended periods. The military has also been the major sponsor of research and development in the robotics and machine vision areas of AI. Major projects include autonomous land and aircraft vehicles, automatic target recognition systems, sonar discrimination systems, and navigation aids.

AI systems proved their strategic value in support of operations Desert Shield and Desert Storm. For example, *DART* (Dynamic Analysis and Replanning Tool) solved the logistical nightmare of moving the U.S. military assets to the Saudi Desert. The application was developed to schedule the transportation of all U.S. personnel and materials such as vehicles, food, and ammunition from Europe to Saudi Arabia. This one application alone reportedly more than offset all the money the Advanced Research Projects Agency had funneled into AI research in the last 30 years.

Another example is *AALPS* (Automated Airload Planning System), a military airlift load planner used by the Army and Air Force to maintain the aircraft's center of gravity, through evaluation of the shape and weight of each piece of cargo. *AALPS* was designed using a graphical interface depicting the aircraft, and point and click-on icon representations of helicopters, trucks, and other cargo are used to position cargo in the aircraft hold. *AALPS* reduced the time required to generate and modify cargo loads from about a week to about one hour.

AI is also used to monitor the Nuclear Test Ban Treaty through an intelligent system, IMS (Intelligent Monitoring System), that automatically detects, locates and identifies underground nuclear tests. It incorporates expert systems, fuzzy logic, neural networks and semi-automated knowledge acquisition.

Applications Component - Commercial - The profile of use in the private sector is different from the military. Where the military community develops a very large and expensive AI system from beginning to end to meet a particular objective, the commercial sector uses more off the shelf, less expensive AI systems to help them achieve more micro-efficiency objectives. AI technology is used in virtually every sector of the commercial economy. About 70-80 percent of the Fortune 500 firms now use AI to varying degrees. Major application areas are in manufacturing, diagnostics, tutoring, financial services, transportation, and data management.

The most applications (about 25 percent) were reported in the manufacturing sector. These included applications in the chemical, steel, auto, electronics, computer, aerospace, and plastics industries. They involved design and engineering, process control, scheduling and planning, part making, factory automation, inspection, and monitoring. AI is also a core technology used in computer integrated manufacturing. In the data management area, for example, an AI program automatically processes and indexes newswires into almost 700 categories for Reuters News Service. For this system alone, savings were estimated at more than $1.25 million in a recent year.

Another major commercial category is diagnostics and testing. Diagnostic systems are used to examine aircraft engines, human hearing, telephone networks, manufacturing machinery and other types of equipment, energy pipelines, ground water and hazardous materials.

A third major segment is transportation services. AI is utilized for traffic management systems, aircraft maintenance operations, airport gate scheduling, railroad planning and forecasting and barge to tow boat assignments. These AI systems are used to manage and draw optimal scheduling decisions from large volume, complex and dynamic databases that would overwhelm human beings.

AI has improved the competitiveness of its commercial users by increasing productivity, improving quality, and augmenting marketing. It has also expanded user capabilities and task performance in ways not previously feasible. It can also induce faster, more consistent and accurate communications, improve service, and improve such intangibles as company or organizational images and customer satisfaction.

AI Technology: A Long Range Perspective

AI technology is a small, but influential part of the much larger general software industry. It represents the leading edge of computer software. AI research undertaken years ago has contributed a great deal to the usefulness of today's conventional software. For example, AI research in the 1960s and 1970s was directly responsible for windowing, spreadsheets, e-mail, spell-checkers, chess programs, and many other system components.

AI is still an emerging technology. Continued research is essential to its long-term development. While many AI techniques have attained commercial viability, improvements are needed to further expand markets. In other cases, such as machine learning and robotics, major research remains undone.

Three major points need to be understood about AI as an emerging technology. First, it is revolutionary in that it potentially raises productivity 10-fold or more, and requires many changes in the mind-set of management and people using it. It can, therefore, have an enormous impact on international competitiveness. Second, as an emerging technology, AI requires sustained long-term research. However, AI lacks a solid constituency (critical mass) in the commercial market to support that research. Federal leadership and patronage, critical to AI's long-term development, is declining primarily due to DOD spending cuts.

Third, AI is experiencing market acceptance problems. It can take a major education effort to introduce a new technology, which from a vendor's standpoint means extensive customer consultations, and salesmanship. From a customer's standpoint it means investment in worker training and equipment. If the AI technology is revolutionary, it may even entail a re-engineering of entire business organizations, and require an often unwelcome paradigm shift in the thinking of business management.

Lack of market acceptance is also partly related to the fact that AI technology came out of the laboratory before the corporate computing world was ready. In 1980, corporate computing was mainframe based, centrally controlled, and focused on accounting and payroll. Early AI vendors tried to get corporate computing management to re-engineer the way they used these systems before corporate managers understood the dramatic changes taking place in computer technology. However, advances in computer technology have forced change. Now corporate computing is becoming multi-platformed, decentralized, and involved in every facet of the business. This change is favorable to the use of AI systems and the other components of the knowledge automation industry.

Government Role

As previously noted, the Federal Government, notably DoD, has played the preeminent role in the development and commercialization of AI technology. In the future, as AI technology becomes more a part of mainstream software and simpler to use, non-Defense agencies can be expected to increase their use of AI relative to Defense usage. However, Defense will remain the major user for the foreseeable future. Perhaps more importantly, Defense will continue taking AI technology out of the lab and fielding first time applications.

The government's role can be divided generally into four elements --

1) **Fund Research**: Sponsor funding for basic and applied research.

2) **Purchase the Product**: Develop and deploy new and existing AI techniques that a) enable the government to accomplish its mission; b) improve government efficiency and services to the public; c) demonstrate the feasibility of the technology; and d) provide an initial market for the private sector.

3) **Manage Business Environment**: Provide legal, regulatory, and educational infrastructure to foster the development and use of AI technologies.

4) **Leverage Risk**: Provide leadership in promoting cooperative agreements. Promote dual use and technology transfer to private sector and stimulate private investments in new technologies.

Government Funded Research - Universities look to the Federal Government for about 80 percent of their research funding. The surveyed universities were in general agreement that the Government needs to continue funding basic and applied AI research if the United States is to maintain its leading position. AI companies (about 75 percent of funding from USG) also look for the Government to fund research. The consensus viewpoint of companies surveyed is that the commercial sector will not pay for the required AI research, particularly at the basic level; therefore, they look to the Government to fill that role.

Government AI Purchases - Statistics on the use of AI in the Government exist only on an partial basis. Comprehensive data are not available, and no system is in place to collect it. However, the Government is unmistakably the largest single user of the technology. In so doing, the Government has contributed steadily to AI's development as a commercial product. On numerous occasions DoD has been the first user of new AI technology, demonstrating its feasibility.

Many government agencies, or bureaus within those agencies, have formalized AI groups that understand the technology and actively promote the technology transfer of AI within their areas. The U.S. Army has the largest such group by far. The Army group, with several hundred people, has many ongoing AI projects. One particularly impressive AI project is called *Blacksmith*. *Blacksmith*, which manages over a terabyte (trillion bytes) of data, is a management decision tool that will enable military planners to simulate with high precision how a change in policy, however large or small, will effect Army force structure and capabilities throughout the world, and thus allow decisions to be made with foreknowledge of the consequences.

The Navy and Air Force were slower to adopt the concept, but they now too have such AI groups. Other agencies, such as NASA, the Internal Revenue Service, Social Security Administration, National Library of Medicine, and FBI, to name a few have also adopted the specialized group approach.

Defense Cutbacks

The DoD has long been the major patron of AI research, and the leading developer and user of AI in the world. A decline in the Department's AI activities could severely slow down basic and applied research as well as AI development. Such a slowdown would adversely affect the long-term competitiveness of the United States. A major concern is that as the Defense budget declines, other Federal agencies will not take up the slack.

University AI labs and private think tanks will absorb the brunt of Defense cutbacks in AI research. A major AI institution characterized federally funded research as a national resource in terms of transfer of government supported technology developments to the commercial sector. Projected downturns in DoD R&D expenditures will impair technology transfer efforts in AI and many other fields and, as a result, impact U.S. competitiveness.

Among AI companies, the most common impact of declining defense funding will be registered as declines in sales to Defense (and to prime contractors). Many AI vendors noted that defense sales have already declined, some sharply. Other companies reported a negative impact on their research activities, and on jobs. Some firms also reported the impact of Defense cuts would be minimal, but added that there could be increased competition as firms shift from defense to commercial sectors.

International Standing

In 1993 the United States AI market was about twice the size of the rest of the world combined. However, experts agree that this huge lead is diminishing as other countries appear to be incorporating AI systems at a faster rate than is occurring in the United States. Technology application appears to be a key factor in the competitive world market. In this respect, other nations, with stronger collaborative efforts between government, industry, and academia, are doing a better job in applying AI technology than the United States.

SUMMARY OF COMPANY AND UNIVERSITY VIEWS OF U.S. INTERNATIONAL STANDING IN ARTIFICIAL INTELLIGENCE

U.S. Standing	Companies	% of Total	Universities	% of Total
U.S. leading	3	8.8%	2	9.1%
U.S. leading, but losing ground in some areas	4	11.8%	1	4.5%
U.S. leading, but losing ground in most areas	7	20.6%	3	13.1%
U.S. lead eroding, behind in some areas	17	50.0%	13	59.1%
U.S. lead eroding, behind in more and more areas	3	8.8%	3	13.6%
Total Responses	34	100.0%	22	100.0%

Universities responding to the OIRA survey noted that the United States is beginning to fall or has already fallen behind international competitors, particularly Japan, in applications of AI. This applies to consumer products, the ability to integrate expert systems with conventional systems, the development of very-large knowledge bases, and business-sector investment in knowledge-based technology. Also, the United States is behind in the design of intelligent computers which can "reason" over large volumes of data. In addition, the United States lacks database, image, and software standardization, which is having a mixed effect; it encourages versatility but discourages large cooperative projects.

Universities also stated that the United States is behind Japanese production of fuzzy logic hardware and software because of Japan's industrial policy, investment environment, and pragmatic acceptance of unconventional technology. Japan has an "implementation based" research community. Japan took the lead in applying fuzzy systems theory to develop intelligent control systems, and is advancing rapidly with industrial and commercial applications. American institutions need to continue taking the lead in theory, but at the same time develop new ways for turning theory into applications. According to the university respondents, the United States is also beginning to fall behind Germany is some areas of knowledge representation and reasoning.

U.S. AI firms responding to the OIRA survey had similar comments on overall U.S. competitiveness. One company stated that, "In the United States money is spent on research that sits on the shelf." Another company said, "We are falling behind because our government does not support the transfer of R&D technology applications. This is especially serious with decreasing defense spending." Some firms noted that the United States is falling behind in the use of AI by industry. Although we still lead in many manufacturing sectors, the industry's reluctance to use new technology products will allow international competitors to catch up with and overtake the United States.

Various companies reported that the United States is losing ground or is now behind in certain aspects of robotics, neural networks, expert systems, fuzzy logic, and machine learning. In machine learning, most of the key people in inductive logic programming and first order concepts of learning are in Europe and Australia, not in the United States.

Overall, U.S. universities and firms are in general agreement that the United States is: 1) behind Japan in nearly all aspects of fuzzy logic; 2) losing ground in generally all categories of AI research; 3) behind Europe in establishing consortia; and 4) significantly behind the rest of the world in commercializing AI technologies.

The Department of Commerce's Office of Foreign Availability (OFA) in the Bureau of Export Administration provided a section for this assessment on AI activities in foreign countries. The OFA report documented the partnerships between government, industry, and academia in Europe and Japan that provide the mechanisms for successfully commercializing new AI technology. For example, the Europeans instituted the European Strategic Program for Research and Development in Information Technology (ESPRIT) under which European Common Market nations collectively sponsor research partnerships in AI and related technology. In Japan, MITI recently established the 10-year Real World Computing (RWC) program that will focus on optical computing, massively parallel processing, and neural systems.

Conclusions

1. AI is an emerging technology of strategic importance to the military and of increasing importance to the international competitiveness of U.S. corporations.

2. The U.S. leads the world in nearly all aspects of AI technology, largely due to over 30 years of patronage by the Department of Defense.

3. AI systems, large and small, often result in a ten-fold or more productivity increase. These increases are possible because knowledge (i.e., the ability to take a specific action to

achieve a goal on given information) is the most underutilized asset in any organization. The best AI systems save companies (and the government) millions of dollars a year.

4. Cutting the Defense budget has resulted in a smaller share of the research dollar going to basic research for AI. It has also resulted in cuts in total research that will not be made up by increases from other Federal agencies. University AI labs and private think tanks will be primarily impacted by these cuts.

5. The commercial sector will not support adequate basic AI research nor form AI consortia without Federal involvement. Declining R&D spending will have a negative effect on AI technology developments and, over the longer term, U.S. competitiveness.

6. The slow rate of AI commercialization appears to be a weakness of America's AI Industry. Japan and Western Europe, with stronger collaborative efforts between government, industry, and academia appear to be commercializing and incorporating AI systems at a faster rate then is occurring in the U.S.. This is diminishing the U.S. lead in market share.

7. Corporate computing is shifting from centrally controlled mainframe based to widely distributed multi-platform based. Knowledge automation of widely scattered organization information is growing in importance and is stimulating the market for AI technology.

8. Statistical tracking of AI is inadequate to develop informed policy options.

Recommendations

- o Federal leadership is needed in establishing an Artificial Intelligence Experts Group, with industry, academic and government participation (representing commercial and military interests) to provide stability, coordination, focus, and direction to such issues as AI research and development funding, commercialization/technology application, and software development productivity.

- o The AI Experts Group could also stimulate the formation of consortia and regional labs to research, develop, and deploy AI technologies as well as address the issue of finding and training qualified AI workers, an issue that was raised by many AI companies and universities.

- o The Federal Government needs to strengthen its statistical tracking of AI technology research and development expenditures and purchases of AI systems by both the public and private sectors.

CRITICAL TECHNOLOGY ASSESSMENT OF THE U.S. ARTIFICIAL INTELLIGENCE SECTOR

1. BACKGROUND

This critical technology assessment of U.S. artificial intelligence (AI) was initiated under Section 825 of the Defense Authorization Act for Fiscal Year 1991. This section of the law requires the Secretary of Defense (acting through the Under Secretary for Acquisition) and the Secretary of Commerce (acting through the Under Secretary for Export Administration) to submit annual reports to the Armed Services Committees of the Senate and the House of Representatives on the financial status and production base of industries supporting technologies deemed by the Department of Defense (DoD) as critical to the performance of current and next generation weapon systems. The National Defense Authorization Act of Fiscal Year 1993, Section 4215, further expands the scope and requirement for technology and defense industrial base capability assessments.

The primary objective of these assessments is to provide industry executives and government policy makers with comprehensive information and analysis on the production and technology status, economic performance, and international competitiveness of private sector firms involved in critical technologies, in light of declining defense budgets. While DoD has deemed these technologies essential to the development of the next generation of weapon systems, they are also crucial to the nation's ability to compete in the global economy. Not surprisingly, almost all of the DoD critical technologies are also found on the Department of Commerce's 1990 list of Emerging Technologies and the Office of the Science and Technology Policy's 1991 list of National Critical Technologies.

Six of the DoD critical technologies were initially selected for review and submission to the Congress. Artificial intelligence is one of the six chosen; the other assessments cover Advanced Ceramics, Advanced Composites, Flexible Computer Integrated Manufacturing (later withdrawn), Optoelectronics, and Superconductivity.

The Department of Commerce's Office of Industrial Resource Administration (OIRA), Strategic Analysis Division, is the unit within the Bureau of Export Administration that is responsible for conducting these critical technology assessments. For each technology OIRA created an advisory team whose members were drawn from the Department of Commerce's International Trade Administration and Technology Administration (including the National Institute of Standards and Technology - NIST), the Department of Defense's Office of the

Secretary of Defense (OSD) Production Resources Support Office (PRSO), and the White House Office of Science and Technology Policy (OSTP).

Special assistance for this assessment was provided by the Defense Department's Advanced Research Project Agency, the AI groups of the U.S. Air Force, Navy, and Army, the National Science Foundation, and the National Aeronautics and Space Administration. Each of these agencies submitted reports to Commerce on their AI activities. Other agencies from whom information was obtained included the Departments of Treasury, Energy, Agriculture, Labor, Transportation, and Justice, and the Social Security Administration, the Environmental Protection Agency, and the National Archives.

In addition to the public sector, OIRA received assistance from associations, universities, businesses, and expert individuals in the private sector. These entities provided support in the area of survey design and field testing, technical advice, mailing lists, on-site visits, and in establishing company contacts. Professor Bob Smith, Director of the Knowledge Engineering Management Lab of California State University at Long Beach, was particularly helpful and generous with his time. Dr. Roger Knaus, President of Instant Recall, Inc., and Dr. Jon Sticklen, Director of the Intelligent Systems Lab at Michigan State University, both provided many helpful comments.

In accordance with the requirements of the FY 91 and FY 93 National Defense Authorization Acts, the following factors were addressed in each of the critical technology assessments:

A. *The financial ability of U.S. industries supporting these critical technologies:*
1) to conduct research and development relating to critical defense technologies;
2) to apply those technologies to the production of goods and services;
3) to maintain a viable production base in critical areas of defense production and technology in the wake of reductions or terminations in defense procurement; and,
4) to expand the defense production base in national security emergencies.

B. *Additional analysis was undertaken on such factors as:*
1) trends in profitability, investment, research and development, and debt burden of businesses involved in research on, development of, and application of critical defense technologies;
2) international competitiveness and market trends;

3) consequences of mergers, acquisitions and takeovers of such businesses;
4) effects of dependence on foreign or foreign-owned suppliers;
5) results of Defense spending for critical technologies in the current fiscal year, as well as the likely future levels;
6) efforts of Defense to expand the use of commercial technology and equipment; and,
7) the need and efforts of industry in the area of defense conversion.

With industry and interagency assistance, OIRA devised a comprehensive questionnaire to collect information to respond to the assessment factors listed above. The questionnaire was field tested with regard to availability of data, technical accuracy, clarity of instructions, disclosure and reporting format. As part of this effort, OIRA co-sponsored a Critical Technologies Workshop with NIST on February 6, 1992, to gather and incorporate industry input into our draft survey instruments and assessment outlines for each of the studies. More than 500 representatives from academia, industry, and government attended the workshop with many providing comments on our draft survey forms. About 50 representatives were present at the afternoon session devoted to the AI assessment.

OIRA disseminated the separate questionnaires to U.S. industry, and selected U.S. Government laboratories and universities under authority of the Defense Production Act of 1950 (DPA), as amended, and related Executive Order 12656. Section 705 of the DPA authorizes the Department of Commerce to collect information when necessary to accomplish analytical activities regarding the domestic defense industrial base.

To enhance Commerce's effort to assess the industry's international competitiveness, BXA's Office of Foreign Availability (OFA) conducted a review of the efforts of leading foreign companies, governments, and research institutions in the technologies. To conduct this review OFA contacted industry specialists in leading domestic and foreign firms, as well as in government agencies and universities. Department of Commerce foreign commercial officers in U.S. embassies and consulates in Europe and Asia also collected and forwarded information to OFA to supplement the data collected from industry. A summary of OFA's review is included in the international portion of this report.

2. INTRODUCTION

The computer is proving to be the most versatile and potentially useful machine ever invented. Recent advances in computer hardware are now making software technologies such as artificial intelligence (AI) more important and accessible to everyone.[1] The real value of the computer is the smartness and friendliness of the software used to run it. Software products are transforming the computer from "number crunching" idiot savants into strategic assets that today are determining the competitive viability of our largest corporations. AI is at the heart of this rapidly unfolding drama and will have much to do with the shape of the next century.

The past 40 years has seen computer technology evolve through three major stages. The 1950s marked the beginning - the era of data processing. Numerical and factual data could be collected, processed, and validated at speeds never achieved before. The 1960s and 1970s witnessed the emergence of information management. Information and data could now be structured, organized, and manipulated to support and accelerate business processes.

The 1980s marked the beginning of the "knowledge" era. Computer languages and tools were developed which could capture dynamic kinds of information and human expertise. This enabled expert or knowledge-based systems to be built and used commercially which could actually solve complex tasks more quickly and accurately than human beings.

> **The challenge of the 1990s and beyond is to extend and integrate knowledge, information and data management technology to the point where problems of any level can be solved successfully by having access to the data, information and knowledge in an automated system.**[2]

---

[1]Many observers believe the computer is actually behind such major events as the collapse of the Soviet Union and the downsizing of large corporations. This belief is based on the understanding that centralized hierarchies are inimical to the free flow of information in a world where the time to respond to breaking information is growing ever shorter.

[2]Taken from a pamphlet prepared by Carnegie Group, Inc. in 1990, on their "Initiative for Managing Knowledge Assets" (IMKA). IMKA was a consortium of Digital Equipment Corporation, Ford Motor Company, Texas Instruments, US WEST, and the Carnegie Group formed to develop software technology to capture and manage complex, corporate-wide knowledge and turn it into a strategic asset for the companies involved.

Our society has become dependent on "knowledge" for its survival. How to make things, deal with complexity, and raise and educate our children are at the core of modern life, business, and culture. According to Joe Carter, a partner with the consulting firm of Andersen Consulting, Inc., the knowledge of how manufactured goods are built and how they work makes up 70 percent of their development costs. And in service businesses, such as selling mutual funds, that percentage is about 90 percent. Yet, Carter asserts, "knowledge is the most under used asset in any organization." If you can somehow transfer knowledge - experience and know-how - from human brains into computer programs, you can leverage knowledge assets to the hilt.[3]

2.1 Artificial Intelligence: What is it?

The term *Artificial Intelligence* (AI) was coined in 1956 by John McCarthy, then an assistant mathematics professor at Dartmouth University, at a meeting among pioneering scientists of the field in Hanover, New Hampshire known as the Dartmouth Conference.

> **In simple terms, AI refers to highly engineered computer software programs used to make computers do things that appear intelligent - such as reason, learn, create, understand human speech, or solve problems.**

AI is both a science and technology. It is the science that investigates intelligence and its replication in a machine. AI has borrowed from other scientific disciplines to evolve into a multi-disciplined approach to understanding intelligence. Other sciences that contribute to the field includes cognitive science, neurology, genetics, psychology, mathematics and statistics, physics, philosophy, linguistics, engineering, and business. As a technology, AI is used to automate the process of extracting or synthesizing knowledge (i.e., knowing what to do) from information and databases.

> **AI has two major goals: 1) to study the nature of intelligence, and 2) to make a computer (or machine) simulate intelligent behavior.**

[3]Business Week, March 2, 1992, p.98

Both of these goals have proven extremely difficult, much more so than originally anticipated. In fact, disagreement persists among AI professionals on a precise definition of artificial intelligence. For example, does it matter if human cognition is the model for AI programming, or is any form of intelligent simulation acceptable? Also, is the goal of AI to build machines that actually are intelligent, or to build machines that simulate intelligence? In large part because of this complexity, and the fact that AI remains a young science, hardened theories have not yet emerged.

AI has evolved into many specialized areas and approaches to the problem. Major branches of AI include:

automatic programming	knowledge representation	planning
decision making	knowledge acquisition	robotics
expert data base systems	logic programming	search
expert systems	machine learning	speech recognition
fuzzy logic	natural language processing	theorem proving
game playing	neural networks	uncertainty
general problem solving	pattern matching	understanding systems
intelligent computer-aided instruction	pattern recognition	vision

(See Appendix A for definitions of these areas.)

At the minimum, an AI system must be capable of a "***flexible***" response. In addition, according to some, an AI system must also be able to "***learn***" from experience. AI systems must contain the following characteristics (which demonstrate the functional relationships between many of the areas of AI research):

a) Receive input from its environment.
b) Determine an action or response.
c) Deliver an output back.

A method for interpreting the input is needed. This leads to AI research into *natural language*, *speech understanding*, and *vision*. The interpretation of the input must be represented in some form that can be manipulated by the machine. For this problem, techniques of *knowledge representation* are invoked. The interpretation, together with knowledge obtained previously, is internally manipulated by a mechanism or algorithm to arrive at an internal representation of the response or action. This requires techniques of *expert reasoning*, *common sense reasoning*, *problem solving*, *planning*, *signal interpretation (vision)*, and *learning*. Finally, the system must construct a response that will be effective in its environment. This requires techniques of *natural-language generation*.

2.2 AI Products

Major AI products include expert systems, neural networks, fuzzy logic, and natural language systems. Expert systems, often referred to as knowledge-based systems, have experienced the most success to date. Expert systems are considered a mature technology in the sense they have established themselves in the marketplace. Neural networks are projected to grow at a fast rate in the next decade, while fuzzy logic is receiving much media attention. Recent technical progress in the areas of natural language understanding and speech recognition is expected to result in their commercial growth as well.

As a trend, AI techniques are increasingly being embedded or integrated with other software systems, which makes their use very difficult to monitor. Also, the various types of AI are now more frequently commingled to take advantage of their specialized strengths. Virtually all new software uses AI systems or techniques, or techniques that resulted from AI research.

An expert system is a computer program that uses "expert" knowledge and human-like reasoning or "logic" to reach conclusions. The knowledge is carefully organized into a ***domain*** in such a way that it can be quickly accessed and the "knowledge" retrieved. The organization of the knowledge is known as ***knowledge representation***, one of the most critical areas of AI. A common method to represent knowledge is with "if..then" phrases (or rules), that point toward a conclusion or "goal". A control mechanism, the ***inference engine***, makes logical deductions based on the knowledge base and facts supplied by the user to reach a solution to the user's problem.

A knowledge-based system acts like an expert consultant in predicting the outcomes of events or diagnosing problems. It derives most of its power from its knowledge. Other components of an expert system include a user interface, a knowledge-acquisition module, and an explanatory interface which explains how conclusions were reached. A select few rule-based systems have over 10,000 rules. However, most have a few hundred.

Expert systems were the first AI systems to find their way out of the research lab and enter the commercial world. The first system, called XCON (eXpert CONfigurator), was introduced in 1981 by Digital Equipment (DEC). XCON was developed jointly by John McDermott of Carnegie-Mellon University and DEC to configure DEC's VAX computers at individual customer's sites. A DEC brochure hailed the system as "the world's first expert system used routinely in an industrial environment."

The success of XCON reverberated around the world and contributed to the exponential growth of expert systems during the 1980s. By 1985, about 50 systems were fielded. At the end of the following year another 300 were in place. By 1987, the number grew to 1,100, 2,200 in 1988, and by 1992, reached 12,500.[4] Today, many thousands of systems are used in all fields of business. Expert systems are designed for such diverse applications as assessing credit risks, running factories, and advising mechanics.

A neural network can be implemented either with a program running on a general-purpose computer, or with a special-purpose circuit board or chip. In the chip implementation, transistor circuits have an internal structure imitative of the human brain's interconnected system of neurons. In a neural network, transistor circuits are the electronic versions of neurons, and variable resistors represent the synapses between the neurons. Electric signals received by the transistor circuits are either inhibited or enhanced, depending upon the task the neural network is performing, when passed on to neighboring circuits, in a fashion modeled to the way in which the brain's neurons pass on electrochemical signals.

Neural networks do not follow rigidly programmed rules, as more conventional digital computers do. Rather, they build an information base through a trial-and-error method. A programmer, for example, can digitally input a photographic image for a neural network to identify, and it will "guess" which circuits to "fire" (activate) to identify the photograph and output the correct answer. Pathways between individual circuits are "strengthened" (resistance turned down) when a task is performed correctly and "weakened" (resistance turned up) if performed incorrectly. In this way a neural network "learns" from its mistakes

[4]*Creating Expert Systems For Business and Industry*, 1990, by Paul Harmon and Brian Sawyer, p. 15, and *Expert Systems Catalog of Applications*, 1993, by John Durkin, p. v.

and gives more accurate output with each repetition of a task.

Simple neural networks were built in the late 1950s, but little progress was made in the field until more powerful network designs and parallel processing techniques were developed a few decades later. Current neural networks are still far simpler than the human brain, and many exist only as computer simulations.

Neural networks are good at pattern recognition, and making sense of very "large" sets of unstructured data.

Neural nets are used to analyze the stock market, appraise loan applications, and perform medical analysis. The military uses neural nets to drive prototype unmanned vehicles and for target recognition.

Fuzzy logic is a concept derived from a branch of mathematical theory of fuzzy sets. Unlike classical (Aristotelian) theory that recognizes statements as only "true" or "false", or "1" or "0" as represented in digital computers, fuzzy logic expresses terms such as "maybe false," or "kind of true". Objects are not placed in sets with sharp boundaries as required in classical logic, but form fuzzy sets without precise definition, say "fast cars" or "warm water" or "long waiting period." In general, fuzzy logic, when applied to computers allows them to emulate an approximation of the human reasoning process, quantify imprecise information, and make decisions based on vague and incomplete data. By applying a "defuzzification" process, the technique creates definite conclusions.

The fuzzy set theory was introduced by Professor Lofti Zadeh at Berkeley in 1965. In the United States, the AI community concentrated on symbolic representation (traditional AI), where Zadeh's theories did not take hold. The United States used systems of differential equations and other traditional control theory techniques to solve these kinds of problems. While Zadeh's original concept remained a topic of debate in the United States, the Japanese embraced it and rapidly turned the research idea into successful commercial applications. As a result, the Japanese lead the world in this area.

Fuzzy logic is used as a control device in products such as washing machines, elevators, cameras, and anti-lock braking systems. For example, conventional washing machines are set to "inflexible" standard cycles (wash, rinse, spin) without special consideration given to wide variations in load size and other variables. Fuzzy logic provides the kind of flexibility that adjusts machine cycles and water levels based on load size and griminess of the water. While the standard cycle is set long enough to ensure a normal load gets cleaned, fuzzy logic optimizes the operation, and generally saves time, water, and energy.

The business community perceives AI in more cautious terms. In fact, the term and the field have fallen in popularity in recent years, in part because of a history of exaggerated claims and unfilled promises by AI proponents. Vendors of AI systems that were proud to be associated with AI just a few years ago, now avoid the connection. Some computer magazines have gone as far as to declare AI all but dead.

Despite this perception, AI systems continue making powerful and even "revolutionary" contributions in the commercial area, although less conspicuously. If measured against the grand goal (i.e., the "quest for artificial intelligence"), AI programs remain primitive compared to the learning capabilities, intuitive reasoning, and common sense characteristic of the human brain. AI is no longer being projected as the centerpiece of software, but as one of a group of software tools that businesses need to power and integrate computers, and thereby leverage their organizations' productivity.

3. THE STATE OF AI TECHNOLOGY

3.1 AI Now and in the Future

Artificial intelligence has come a long way since its formal birth at the Dartmouth Conference of 1956. (Information on the historical origins and development of AI can be found in Appendix B.) As the field began to develop in the 1970s, AI research centers were formed at increasing numbers of educational institutions. Today, AI courses are offered by many computer science departments. More recent developments in artificial intelligence include voice and graphic pattern recognition software using neural networks. For example, the U.S. Post Office uses a pattern recognition machine to sort letters by zip code. Also, voice-recognition is already in use (in voice menus) by a number of businesses, including the telephone companies. Voice recognition is used to transform speech to written text for the deaf, and character recognition is used to transform the written word to speech for the blind. As these various aspects of "thought" are developed for computer processing, the field of "robotics" takes on new significance.

The current proficiency achieved in expert systems technology is the result of nearly 30 years of research. However, within the past five years, several key advances mark its maturity into a reliable and mature software tool. It is no longer necessary to build an expert system from scratch. A seasoned AI industry has transformed the lessons learned about solving logical problems into structured software shells or skeletons that can be completed by end users. For example, the five-year MYCIN project, an expert system with only 475 rules that took 20 person-years to develop (i.e., two weeks per rule), could now be done in about three hours per rule.[5] This is more than a 100-fold increase in speed; now one person could do it in about two months. In the near term, forecasters predict the appearance of task-oriented shells that will improve the performance of inference engines marketed for specific domains.

Expert systems, regardless of how they are programmed, typically can now run on many different computer platforms - PCs to mainframes. They can also be embedded into standard

[5]*AI: The Tumultuous History of the Search for Artificial Intelligence*, by Daniel Crevier, 1993, p. 200.

applications, and can process information stored in database management systems. The ability to integrate expert systems technology with existing operations has eliminated one of the key obstacles to widespread adoption. The development of alternate methods of knowledge representation, such as case-based and analogical reasoning, opens the technology to a wider range of domains and more complicated problems.

An impediment in deploying expert systems concerns the availability of knowledge bases. Many current expert system tools do not promote structured knowledge base development, and knowledge languages have not been standardized. It is, therefore, difficult to reuse a knowledge base built for one application in another, thus increasing the cost and uncertainty of expert system development.

A fundamental weakness of expert systems is "brittleness". They are useless (they break) when asked to solve a problem or reason about a task that falls outside their knowledge realm. For example, a medical diagnostic system may recognize a 10-year old motor vehicle covered with rust-spots as "measles," and recommend bed, a comic book, and aspirin. The knowledge base does not include "common sense" - here, the observation that inanimate objects do not contract infectious diseases.

Long-term research addressing the more fundamental problem of application brittleness focuses on two areas: 1) the development of very large "basic" knowledge bases, and 2) machine learning. The industry's key approach to overcoming application fragility is to develop very large "basic" or common sense knowledge bases. The most comprehensive project is taking place at the Microelectronics and Computer Technology Corporation (MCC), formed in 1984.

In a collaborative 10-year effort funded by six major computer and communications corporations (and ARPA initially) that began in 1985, MCC is developing an AI program called *CYC* (enCYClopedia). Although there are many skeptics, when finished CYC is supposed to be able to read and understand passages from an encyclopedia. Behind these passages lies a wealth of hidden "common sense knowledge" we all share from life experience.

For example, we all know to "open the door before walking through it," or "if the bird is dead, it can't fly." Like a human, a computer has to have this etched (programmed into a knowledge domain) into its memory, and be able to use it at a moment's notice. This is not an easy task. CYC now has about 115 megabytes (roughly equivalent to all the information a 26 volume encyclopedia). When confronted by this issue, some AI scientists experience "physics envy," i.e., the ability to say "all encompassing" things with short phrases or formulas - "$E=MC^2$."

Another limit (or perspective) on AI as currently evolved is its inability to codify the extremely complex qualities of human insight, understanding, and creativity, which give humans a certain flexibility and know-how. These qualities, collectively referred to as "tacit knowledge," endow a person with the ability to bring seemingly unrelated experience(s) to bear on problems and situations." Humans are sentient, curious, learning, intuitive, willful, and motivated; these are survival skills (or coping tools). Recently, theories developed years ago on tacit knowledge have resurfaced.[6]

Tacit knowledge is the power to innovate and improve. Not surprisingly, the collapse of the Soviet Union, a centrally controlled organizational structure which stifled the expression of tacit knowledge, reinvigorated these theories. This event and others, such as the downsizing of corporate bureaucracies, reinforced the belief that humans function more effectively in decentralized environments where they are both empowered and given responsibility to think for themselves. When focused on economic matters, human expertise and flexibility are indeed a potent force.

Lee Bloomquist, Principal Engineer in Research at Steelcase Incorporated, has the view that tacit knowledge ultimately determines national competitiveness. For example, if two factories are located in different places and given equal access to the very same labor markets and the very same technology markets; he contends the firm that best manages its people's tacit knowledge will achieve a sustainable competitive advantage. The same holds for advantage among nations.

Many AI scientists have gained a new appreciation for human "smartness" by unsuccessfully

---

[6]Frederick von Hayek, a noted economist and philosopher of science, claimed in the 1940s that there would always exist within the market what we today call "tacit knowledge." As von Hayek saw it, tacit knowledge is that knowledge which can never be codified into information for use by a centralized hierarchy and effectively be used to direct an economy in all its details. He postulated that "the crucial knowledge by which an economy thrives is inextricably embedded within the market itself - it takes the form of tacit knowledge." With the economic failure of the Soviet Union, von Hayek's ideas are getting renewed attention.

trying to codify it into computer programs. AI, on the other hand, is extremely useful in codifying existing knowledge and know-how. Throughout recorded history, tacit knowledge has generated codifiable knowledge. For example, the airplane was "imagined" long before it became a reality. However, the reality is codifiable. The car was invented by a human who "creatively" rearranged and assembled it from existing components. The car is codifiable; it can be described in every detail. Factory automation, chess games, the pyramids, and shoelaces are all codifiable, but their common origin is tacit knowledge. AI's power is to capture expertise and distribute it throughout an organization. In this fashion, AI raises the level of expertise in an organization by distributing it to whomever or wherever it is needed.[7]

Perhaps in the future computers will actually "learn" and etch their own memories with common sense and tacit knowledge. Herbert Dreyfus, a Professor of Philosophy at UC\ Berkeley, has been one of the most articulate skeptics in this area. He points out that the human "body" is a critical part of our learning experience (i.e., it tastes, smells, feels a hot stove, etc.). He believes a computer cannot be programmed to experience these things. So far he is correct.

AI systems can today:[8]

- Help organizations manage "knowledge" assets.
- Help organizations manage complexity.
- Help experts to solve difficult analysis problems.

---

[7]For example, Lee Bloomquist noted that AI technology can be used to capture the expertise of "success" stories - such as best manufacturing practices, or TQM (Total Quality Management) - and present them interactively and individually to raise the level of understanding of people throughout an organization.

[8]Most of statements about AI now and in the future, with minor modification, came from *Artificial Intelligence*, 3rd Edition, by Patrick Winston, 1992, pp. 6-11. The last two statements under long-term prospects were provided by Dr. Roger Knaus, President of Instant Recall, Inc. As a word of caution, the statements are examples of possibilities more so than predictions, although based on expert opinion. If the past is any guide to the influence of new technologies, many future effects of AI on lifestyle, business, and society are currently unanticipated.

- Help experts to design new devices.

- Learn from examples.

- Provide answers to English (or human language) questions using both structured data and free text.

In the near term, AI system possibilities include:

- In business, computers may help us to locate pertinent information, to schedule work, to allocate resources, and to discover salient regularities in databases.

- In engineering, computers may help us to develop more effective control strategies, to create better designs, to explain past decisions, and to identify future risks.

Long-term prospects include:

- In farming, computer-controlled robots could be used to control pests, prune trees, and selectively harvest mixed crops.

- In manufacturing, computer-controlled robots could do the dangerous and boring assembly, inspection, and maintenance jobs.

- In medical care, computers could help practitioners with diagnosis, monitor patients' conditions, manage treatment, and make beds.

- In household work, computers could give advice on cooking and shopping, clean the floors, mow the lawn, do the laundry, and perform maintenance chores.

- In schools, computers should be able to understand why their students make mistakes, not just react to errors. Computers should act as superbooks, displaying planetary orbits and playing musical scores, thus helping students to understand physics and music.

- o In research and development, computers may search for and organize information from vast world-wide computerized reference libraries.

- o In information processing, computers may assist with more of the detailed design, programming, requirements analysis, knowledge acquisition, verification, and validation

As a final point (or perspective), AI and other software are converging. In the future at some point the distinction between the two may be somewhat meaningless. As AI techniques continue falling into standard and widespread use, software, in general, exhibits more intelligent behavior. As the leading edge of software, AI is an important focus of research that greatly affects all software. If viewed in this broader context, software, in general, and AI, in particular, are both the driving force and the bottleneck to the tremendous potential of the computer. By some estimates software has now ballooned to a $300 billion global industry that supports additional trillions of dollars in commerce. It is at the core of nearly everything done today - from library science, aircraft design, and weather forecasting to credit approval, medical research, and factory automation.

The demand for software (especially customized software made for in-house use by corporations and governments) far exceeds our ability to supply it. A major problem is that most software programs initiated are never deployed because of too many glitches and other deficiencies. Also, the productivity of software development has hardly improved at all in the last 20 years while the complexity of programs has increased. Automated techniques to build software are needed. As one of its most important potentials, AI may find use as the "machine tools" that automate (and revolutionize) the process of software development. In so doing, AI can also help improve the quality, shorten development time, and lower the cost of software development.

3.2 AI Technology Applications

The OIRA survey questionnaire used to gather information from the private sector for this assessment included a question regarding AI system "applications". (See Appendix C - the Survey, Part II, # 2, p. 3). The AI companies were asked to identify three "defense" and three "commercial" applications for their AI products. Three universities also provided this information voluntarily. In the defense market, 35 companies and the three universities reported 80 applications. In the commercial area, forty-three companies (and the universities) reported 105 applications. The total of 185 applications were grouped into general application areas as shown on Table 1.

Table 1.
Defense and Commercial Application Areas of
Artificial Intelligence Technologies
(number reported by area)

Application Areas	Defense	Commercial	Total
Manufacturing	4	25	29
Diagnostics and Testing	15	12	27
Tutoring, Help Aids	12	11	23
Robotics, Vision, Other Sensory	15	5	20
Data/Information Analysis	7	8	15
Transportation Services		11	11
Financial	1	9	10
Planning	8	1	9
Management Support	4	4	6
System Management	0	4	6
Monitoring	7	1	8
Government (Non-DoD)		3	3
Sub-Total	73	94	167
Unspecified, Generic	7	11	18
Grand Total	80	105	185

Source: U.S. Department of Commerce, BXA/OIRA Sector Survey

Nearly all AI applications cited by survey respondents offer some form of decision support at various levels in any given organization. Exceptions would be the teaching aids that were grouped into the tutoring-help aids area, although some of these are also decision aids. The applications use a variety of AI technologies (expert systems, fuzzy logic, neural networks, vision systems, natural language) either alone are in combination, as well as other software technologies (graphics, statistical programs, word processing), to achieve their ends. The applications almost always dramatically increase productivity when applied in appropriate areas. The cost and complexity of the AI systems varies over a wide range, from very cheap to hundreds of thousands of dollars. Some AI programs are newly developed, and advancing AI technology in the process.

Grouping the numerous AI programs into a particular application area is not an exact science. For example, some of the diagnostic and robotic programs are actually applied in the manufacturing area. However, their primary characteristic as diagnostic or robotic led to their separate groupings.

As Table 1 highlights, the areas of heaviest concentration of AI usage in the military market are diagnostics and testing and robotics (each about 19 percent), and tutoring, help aids (15 percent). In the commercial arena, manufacturing (at close to 24 percent) was mentioned more than twice as much as any other category. Diagnostics (at 11.4 percent) was the distant runner-up, with tutoring and transportation services each indicated about 10.5 percent of the time.

3.2.1 Military Applications

Diagnostics and Testing: Expert diagnostic systems are used extensively by the military to troubleshoot and maintain military equipment. An important value of these systems is facilitating repairs by non-experts on equipment that otherwise could be out of commission for extended periods. Expert diagnostic systems also conserve space and weight, and greatly reduce the burden of hauling numerous repair manuals around, which may stand over 10 feet if stacked. In fact, this aspect alone usually more than compensates for their cost.

Moreover, expert diagnostic systems are considerably easier to understand and faster in use than manuals. Some of the expert diagnostic systems in use include the M1-A1 Abrams tank engine, Apache helicopter, air defense Hawk missile, the MX missile inertial measurement unit (for guidance control), the flight test and range safety systems of the F-15 and F-16 fighter aircraft, and several communication and ground support systems.

AI programs are also developed for testing and inspecting equipment or equipment parts. These AI programs, which have been lumped in with the diagnostic family, are used to inspect such items as gas turbine engine blade surfaces using fluorescent penetrant and the shuttle's rocket engine face plates between flights. These AI systems not only speed up the process, but provide more consistent and accurate results, and increase safety at reduced cost.

Robotics: The military has been a major sponsor of research and development in the robotics and machine vision areas of artificial intelligence. Major projects have focused on autonomous land and aircraft vehicles, automatic target recognition systems, sonar discrimination systems, and navigation aids. The U.S. Army is using a robot to patrol material storage areas for indoor security. NASA has used a robot for testing video navigation concepts. Target recognition systems are being developed to find tanks, aircraft, and other objects. One firm reported it is building a "vision library" for enabling a computer to recognize various military images. Firms reported using neural networks to classify (sonar) signals and to segment area reconnaissance imagery.

Tutoring, Help Aids: In the tutoring and help aids area, a university reported development under a military contract of spoken language interfaces with computers. Much progress has been made in this area by a number of different schools and companies. This type of interface should be fairly common in the near future. Until now, the primary person/machine interaction has been by keyboard and mouse. In another area, computer-aided instruction and training is becoming more common and interactive. One tutorial that simulates a Navy ship engine room is used to train U.S. Navy personnel (machinist mates). Help desks are used for training, resolving problems, and accelerating response time to specific questions. For example, an intelligence agency uses a help desk that allows prompt answers to questions related to computers and computer networks. These kinds of AI programs are also useful as help aids for weaving through complex procedures and completing forms by quickly focusing the user on relevant information, and answering specific questions.

More exotic aids include the "Pilot's Associate." Funded by the Advanced Research Projects Agency and the U.S. Air Force, this program is a pilot's cognitive aid to improve situation assessment and mission effectiveness of high performance fighter aircraft. Another application is used by the U.S. Army. The program, called "the day/night adverse weather piloting system," is a pilot's cognitive aid to improve mission effectiveness and survivability of Army helicopters.

Planning: Eight reports were received for planning applications, primarily in the logistics and battle management areas. Each of the Services was mentioned as an end user. One firm provided a knowledge base logistics planning shell to the Army (G4 Staff Planners) for inventory, distribution, and transportation planning. The program captures a given battle plan, analyzes requirements and capabilities, and then advises on logistics supportability. In another instance, the Air Force used a warplanning occupation requirements AI program during Desert Shield and Desert Storm that gathered information on the nature of the conflict and assessed which job categories needed to be activated. It further determined how individuals in one job category could be transistioned to other categories based on the state of the conflict. Also, the U.S. Navy has an expert system that assists with force requirements and the management of asset (ship) allocation and replacement. Other programs are used to generate bills of material, optimize weight distribution of cargo loads in military aircraft, and plan the orderly off-loading of material at a given destination.

Monitoring: ARPA funded development of AI techniques to help monitor the Nuclear Test Ban Treaty. This was undertaken by ARPA's Nuclear Monitoring Research Office, largely through the Science Applications International Corporation (SAIC). The Intelligent Monitoring System (IMS) was developed under this program. IMS is an intelligent system that automatically detects, locates, and identifies underground nuclear tests. It incorporates expert systems, fuzzy logic, neural networks, and semi-automated knowledge acquisition. In another initiative, the U.S. Air Force Technical Applications Center (AFTAC) contracted Sandia National Laboratory in New Mexico to develop, integrate, and install a complete system to monitor global seismic activity. SAIC provided most of the software for this project. The fruits of this project resulted in the AFTAC Distributed Subsurface Network (ADSN) System, similar to IMS, but used for global monitoring.

Several non-nuclear monitoring AI applications were also mentioned. These include a satellite command and control/telemetry monitoring system, the Savannah River Laboratory Surveys for radioactive contamination, an electrical load management and power distribution system, and expert systems used to monitor and control a telecommunications network.

Data/Information Analysis: The military generates huge quantities of raw data in daily activities. AI techniques are increasingly being used to extract information from this data, and in turn, knowledge from the information. AI extraction techniques such as data fusion (combining data from multiple sources) and data mining (searching for patterns) are utilized. For example, in one application an intelligence agency fuses data from multiple sources and derives coherent, non-obvious relationships. Another program uses intelligent software to turn disparate incomplete and uncertain sensor data into representations of knowledge about the real world which is used in combat decision systems. In a different area, the U.S. Coast Guard uses a automatic message processing system to locate ship positions at sea. A Defense agency is using fuzzy logic for threat classification. Another group (unnamed) is providing expert systems for implementation in nuclear debris collection and analysis, and integrated data evaluation in the nuclear plant program. Also, a firm reported its memory-based reasoning AI tool is used by the military for retrieving information from data.

Defense Manufacturing: Defense-related AI manufacturing applications support various weapon systems by speeding, improving, or enabling the manufacture of parts and subcomponents. Those cited by survey respondents are used in the engineering design phase of the manufacturing process. For example, one AI program supports the rapid design of castings and casting molds. Another program captures and organizes design expertise and advises designers on "best practices," and presents prior examples of part and mold designs. A third program supports structural design using neural net-based rapid structural analysis integrated with structural optimization. A fourth application is a shell that supports the manufacture of avionics.

Management Support: This is somewhat of a catch-all category for four applications that did not fit neatly elsewhere. Included are a general decision aid to assist technicians in "fix vs. replace" decisions that balance cost with readiness, and an intelligent combat decision aid/battle management tool. Additionally, a firm reported its AI package provides advanced distributed simulation of battle conditions for semi-automated forces. Another firm reported its parallel neural network package enables a parallel digital computer to emulate a neural network.

3.2.2 Commercial Applications

Manufacturing: As the most commonly mentioned application area (25 times), AI techniques are used widely in the manufacturing sector. Applications were reported in the chemical, steel, auto, electronics, computer, aerospace, and plastics industries. They involve design and engineering, process control, scheduling and planning, part making, factory automation, and monitoring; all of the various AI techniques are used.

For example, a firm is using an AI system to schedule the manufacturing of seat belts and air bags. A paint manufacturer uses an AI color matching system for automotive paint manufacture. Other such color systems are also reported as used in the textile, printing ink, plastics, and food industries. A chemical company uses a process simulator to optimize yields. In another case, a design tool for chemists and chemical engineers "discovers" formulation and processing models, allowing users to provide optimal product properties and processing while minimizing product design time and experimentation. A wax injection machine (used in making casting molds) employs AI consultant software in the production of aircraft engine components. A computer manufacturer uses a real-time, on-line decision making control system for an automated factory. This particular firm reported savings of $12 million a year compared with before AI system results. Another computer manufacturer has deployed a system that enables faster and better processing of plastic injection molded parts.

In other applications, a hybrid neural net/expert system monitors plant operations and alerts the operator that some data or parameter is suspect or out of specification. Also, an advanced control system is being used to control continuous processes and reduce process variations. Another system is scheduling production at a semiconductor manufacturing facility. An auto manufacturer is using an expert system to dispatch automotive assembly. A chemical company is using AI to monitor environmental compliance of one of its plants. AI techniques are also being used to control cement kilns and to design cement slurries for casings. The energy sector utilizes AI to design acid natures for downhole acid jobs. An expert system is used to capture and electronically transfer semiconductor manufacturing knowledge and expertise.

Two AI systems, each called a process advisor, are used in metal production applications, although the use was not specified. However, one is located at a steel coil operation. Another system was at a chemical plant control room monitoring the alarm system, while still another involved the maintenance and operation of chemical process equipment.

Diagnostics and Testing: Diagnostic AI programs are versatile and widely used throughout the economy. Twelve such systems were cited by survey respondents. Most are expert or knowledge based systems that utilize various AI technologies. Diagnostic systems are used to examine aircraft engines, human hearing, telephone networks, manufacturing machinery and other types of equipment, energy pipelines, ground water, and hazardous materials. For example, one versatile program has numerous commercial applications, such as an "intelligent manual" for equipment/process diagnostics, or troubleshooting on integrated circuit making equipment, or start-up, shut-down, operating procedures, and product configuration information. It operates in a user-friendly graphical environment. Another example is an expert system used for diagnosing chronic problems in local phone networks. Also, an expert system is being used to supervise the testing of telecom special service circuits and to provide a diagnosis when problems are found.

Tutoring, Help Aids: Eleven tutoring and help aids applications were reported for the commercial market. Most are used for educational purposes or as help desks to, for instance, speed-up response time by service representatives. In one example, Compaq Computer Corporation is using an AI program called the Support Management Automated Reasoning Technology (SMART). Compaq needed to differentiate itself in the crowded PC market place. In response, it has promised direct phone support for their new lines of computer rather than a dealer only strategy. SMART is a deployed shared knowledge help desk AI application which automates the resolution of problems for customer service representatives. In addition to SMART providing one stop problem resolution, it has helped Compaq gain significant market share.

Another system is HALIPSOS, which was developed by the IBM Paris Scientific Center. Done in Prolog, HALIPSOS creates automatic understanding of French text. Another firm mentioned its "autograder" that uses speech recognition technology to automatically evaluate the pronunciation of foreign language students. An expert system called EXPERTiMENTAL DESIGN helps researchers decide what kind of experimental design is most appropriate for their research. Other applications included automated help desks for computer users and various telephonic types of instruction, response, and training systems.

Transportation Services: Eleven reports of AI applications in the transportation area were received. These included traffic management systems, aircraft maintenance operations, airport gate scheduling, railroad planning and forecasting, and barge to tow boat assignments. These AI systems are used to manage and draw optimal scheduling decisions from large volume, complex, and dynamic databases that would overwhelm human beings. For example, American Airlines has an expert system called MOCA. This is a real time system used in the maintenance operations control center of American Airlines. It supports the maintenance controller in maximizing the utility of the aircraft while complying with all federal and operational maintenance policies. MOCA specifically determines which aircraft will fly which routes. After deploying MOCA, American Airlines improved the utilization of its Boeing 727 fleet by over 10 percent, a very large savings in dollars.

Another air transportation AI program produces an intelligent automated real-time scheduling of aircraft at terminal gates. This system incorporates real time updates of flight schedules to model gate use through a day allowing users to see the downline effects of gate assignment decisions. A second system incorporates real time updates of flight schedules to model flights over a day, taking into account aircraft capabilities and passenger loads. This system allows users to replan under irregular operations such as airport closings, weather, and aircraft breakdown.

In other examples, Burlington Northern and Southern Pacific Transportation Company each use AI techniques to improve operations. Burlington has a system that resolves conflicts in train schedules with track maintenance. Southern Pacific's system assists in projecting railroad car demand by car type and location. Another AI system provides an intelligent automated assignment of barges to tow boats. The system projects boat arrivals at major hubs and generates planning alternatives for breaking apart and building tows. It provides objective analysis of planning alternatives, including comparison of factors like "barge shifting cost," "boat port time," and "last barge profit contribution."

Financial: Nine reports of financial applications were received from survey respondents. In part because of the voluminous data and need for quick response, this is an area where AI techniques, particularly expert systems and neural networks, excelled early. An important system is American Express (AMEX) Company's *Authorizer's Assistant* developed by Inference Corporation. This knowledge based system is one of the best known in the industry. The American Express card is a charge card, not a credit card. In order to manage their business AMEX must perform a full-up credit decision each time the card is used. In simple terms Authorizer's Assistant using AI techniques determines, "is this the person named on the card?" and, "will this person pay if approved?" This system processes tens of thousands of transactions per month and saves American Express millions of dollars a year.

Other programs are used to detect fraud, assess risk, proform financial analysis, manage funds transfer, generate finance applications, and process bank messages. For example, a fraud detection system for demand deposit accounts is used in banks. An AI (object oriented) program using Prolog generates insurance applications. Another program enables the user to build a customized application using expert system technology. For instance, a bank can use the program to model a loan applicant expert system that would indicate, based on criteria, a good/bad risk. Another AI system is used to provide customer support for electronic funds transfer.

Data/Information Analysis: Eight reports were received in this application area. One AI program automatically processes and indexes newswires into almost 700 categories for Reuters News Service. Savings were estimated at more than $1.25 million in a recent year. In another instance, an automated classification tool assigns information in assorted documents to pre-established classes. And another retrieves documents in a textbase searching system. In the medical field an application uses neural nets to identify malignant cells. Other programs are used for statistical analysis, data inquiry and retrieval, and coding or information classification functions.

Robotics, Vision, Other Sensory: While only five reports referencing commercial robotics applications were received, broader use was indicated. For example, one firm reported it had deployed hundreds of commercial machine vision systems, without being specific. Vision technologies, such as pattern recognition, are advancing rapidly. Another firm reported an unspecified number of part feeding, sorting, and inspection type robotic applications. A third firm stated its automated vehicle transports materials in light manufacturing, office buildings, and hospital environments. And a fourth firm noted its robotics were used in the food, pharmaceutical, printing, and traffic related industries. Only one report was more specific - a roving security robot used in commercial buildings and museums.

Most existing robots do not use AI technologies. However, as AI technology (such as image processing) matures, it will become standard in the robotics area. Robots that use AI technologies are distinguished from the more common types by being able to respond "flexibly" to changes in their environment. In contrast, the painting and welding robots commonly used in auto assembly plants are programmed to do a definite sequence of maneuvers. While they can be reprogrammed, all decisions are made in advance.

Other Applications: A wide range of other applications were reported by survey respondents. An AI program, for example, is being used to monitor, diagnose, and foresee emerging problems with a Local Area Network (LAN) system. Expert systems are utilized to control and maintain telecommunication networks, and to automate an electric power grid distribution system. A real time advisory system is assisting the operation of a nuclear power plant. Other systems are helping companies improve customer service and sales support, and make intelligent purchase decisions.

The Federal Government (other than the Defense Department) is also making extensive use of AI techniques. For example, the Department of Commerce, National Oceanic and Atmospheric Administration uses an AI system (THEOPHRASTUS) to predict solar flares. NASA has a system that configures electrical cables for the space shuttle orbiter. The Environmental Protection Agency employs an expert system that automatically sequences treatment for waste water. Although only these three reports were filed, the Federal Government is "the major developer and user" of AI. This is covered more extensively in the section titled "U.S. Government Role in AI Technology."

3.3 How AI Can Increase Competitiveness

> **Intelligent software technology is key to the management and integration of organizational "knowledge assets" that will otherwise be underutilized. The job of management is to know how and where to use AI technology, not whether or not to use it.**

To gain additional insight into the competitiveness issue, the surveyed AI companies were asked to comment on how AI has improved: 1) their own competitiveness; and/or 2) the competitiveness of their customers that use AI systems. If available, they were also asked to provide before and after information that shows cost savings and productivity increases attributable to the implementation of AI software, or the ability to do things not feasible without the use of AI technology. A total of 39 firms responded. Eight of these were large corporations that developed AI systems internally, while all the others were AI vendors.

The company responses were split into three groups as follows: 1) AI vendors that use AI in their internal operations to enhance their own competitiveness (AI Vendors); 2) end-user companies that purchase AI from vendors (End-user Purchasers); and 3) end-users that develop AI in-house (End-user Developers). End-user purchasers were subdivided into two groups. The first group is labeled *general* in that the AI vendor's AI package is used generally, and does not apply to a particular end user. The second group, was labeled *specific* since it applies to a particular end user.

Survey responses were also characterized under four major methods as to how competitiveness was improved. These included: 1) increased productivity; 2) improved quality; 3) augmented marketing; and 4) other methods. The "other methods" category involved either an expansion of user capabilities or the performance of a task not feasible without AI. Some companies cited more than one method in their response. Table 2 summarizes the responses.

Table 2.
AI's Competitive Impact

Number of Companies	Respondent Grouping	Total Methods Cited	Methods of Improving Competitiveness			
			Increases Productivity	Improves Quality	Augments Marketing	Other
12	AI Vendors	14	3	2	9	
10	End-user Purchasers (general)	22	13	5	1	3
9	End-user Purchasers (specific)	12	12			
8	End-user Developers	9	4		1	4
39	Overall Total	57	32	7	11	7

Source: BXA/OIRA AI Sector Survey

The table may be somewhat biased in its emphasis of AI as a contributor to increased productivity, in part due to the wording of the question in the survey. While increased productivity is very important and should not in any way be discounted, it is also paramount to note that AI contributes significantly to user competitiveness in other areas. Perhaps a broader perspective is revealed by the responses of actual AI end users (AI Vendors and

End-user Developers on the table), who cite a variety of factors associated with AI, in addition to increased productivity, that contribute to their competitiveness.

AI technology can break down the rigidity of bureaucracies, and improve management decision making. It can facilitate a more optimal use of resources and create a more flexible and adaptive organization. It can induce faster, more consistent, and accurate communications; improve the quality of products and service; and enhance company or organizational images. Perhaps all this can be measured (mathematically) in terms of improved productivity and a healthier bottom line. However, intangibles such as customer satisfaction, company morale, and expanded sales should not be overlooked.

3.3.1 AI Vendors: AI vendors that use AI techniques in their internal operations (or products) reported it useful as a marketing tool in nine of 14 instances cited. Identification with sophisticated software improved their company image, and promoted good will among their customers. Two vendors stated the demonstration of competence in the use of AI technology was a major influence in obtaining AI contracts, in one case for tens of millions of dollars. One of these companies added it also contributed to non-AI related contracts. Other vendors reported that AI content in their product offerings boosted sales, and is a market plus for value added resellers. One vendor stated, "It is a money maker that increases customer satisfaction and capabilities far beyond its cost."

The vendors also reported significant productivity increases, noting much faster software prototyping and development. One vendor stated AI languages permit a 6-10 fold reduction in development time. AI permits the use of a higher level of abstraction to express the solution to problems.

3.3.2 End-user Purchasers (general): Ten AI vendors cited 22 ways AI increased the competitiveness of end-user purchasers (general). Nearly 60 percent, or 13 of these ways were increases in productivity. For example, one vendor claimed end users can realize a 24 to 1 gain in design applications. Another vendor claimed, that clients can implement their projects up to 10 times faster where AI is applied. A Prolog marketer stated that when software engineers use Prolog for applying logic programming, they improve productivity 5-10 times over other languages such as COBOL and C. A vendor selling a pre-compiler noted his product enables application developers to increase their productivity 2-4 times. Another vendor indicated that many production scheduling applications have actually increased production output more than 100 percent.

Other vendors reported that AI enabled end users to design formulated products more rapidly and more optimally. AI also permits products to be produced less expensively, and reach the market more quickly. And AI's greatest contribution to clients is cost savings, improvement in speed, and productivity enhancement. Another reported that end-users experience reduced software development time and cost, lower maintenance costs, improved operating economics and efficiency, and higher yields. Additionally, both staffing and training time is reduced.

One vendor reported that natural language message processing systems in banks have permitted personnel reductions from decreases in manual processing. A vendor also reported that machine vision systems improved manufacturing efficiency, and another vendor that AI techniques used to configure and layout computers increase productivity. Lastly, AI applications are being built into "re-usable software components" and products that far exceed the cost of development. In many cases, the investment is paid back in less than one year.

Five vendors reported that end users experienced improved quality. One noted its client's products perform better. Another stated its customers have benefitted from improvements in its optical character recognition technology. One end user was reported to have increased sales due to AI use. In the other competitive improvements category, one vendor indicated his clients reported improved safety. Another reported text processing technologies capabilities have provided customers with many significant benefits. In the military arena, one vendor stated military logistics applications provide functionality that would not be possible without the use of advanced AI technologies and techniques.

3.3.3 End-user Purchasers (specific): Nine vendors cited 12 ways AI increased competitiveness in specific applications for particular end users. While this is a small sampling, AI is used throughout the economy in agriculture, manufacturing, transportation, financial services, education, and government.

For example, at a steel mill, AI techniques allowed reduction in staff from four 12-man shifts to four 6-man shifts and lowered the required skill level. In another example, after AI (and other work) was instituted, a metal producer experienced a 92 percent increase in production; possibly millions of dollars were saved on bottom line. A power utility company created an application to monitor and collect information regarding the efficiency of data retrieval. So far, the company has reduced overall systems response time as well as the time spent training their staff. They estimate an annual savings of more than $150,000, a 20 percent productivity gain, and enhanced quality and service.

In the airport gate scheduling area, it now takes one person one day to develop an airport gate scheduling plan that used to take two people an entire week. The airline has successfully managed increased airport traffic and responded more quickly to changes in schedule. In another reported application, intelligent real-time scheduling of aircraft at terminal gates has reduced fuel burn at gate holding areas and improved traveller service and satisfaction.

A car rental company combined the expertise from the sales and pricing groups and created an application that sets rates for car rentals. Because of this work, the company revised corporate policies, streamlined its pricing process, and standardized its price adjustments. The company also gained competitive advantage by maintaining the same work force size while rentals have increased by 25 percent.

A financial and insurance consulting company used AI as an enabler to re-engineer their fixed and variable annuity business. They built a knowledge based system which transformed their organization from one which required state regulation specialists to one where generic processors could perform all work. The benefits of the system were a 70 percent reduction in staff, and a modern efficient business process that streamlined revenue generation and became the lowest unit cost player in their line of business.

A telephone answering service system improved productivity using AI techniques by reducing the length of the average call from 10 to 2 minutes, and increased the volume handled from 2,000 to 13,000 calls. Another company reported an AI tool used to work its computerized reservation system queues at a help desk did the work formerly requiring 20-30 persons. In another example, a similar AI tool monitors a reservation system, reducing idle time and staff, and speeding-up response.

3.3.4 End-user Developers: Eight end-user developers reported 9 instances of how AI improved their competitiveness. Each of these companies is on the Fortune 500 list. While only one mention was directed toward augmenting marketing, it probably speaks for each company in reporting the use of AI technology has "heightened our corporate image and increased product sales."

Three of the companies reported four instances of productivity increases. In one instance, an expert system used in a factory for testing newly assembled computers improved production throughput of the testing process by 20 percent in its first month of operation. A major producer of integrated circuits reported using expert systems and knowledge based systems to help improve manufacturing productivity. This included help in problem detection and diagnostics, as well as help in factory logistics such as scheduling, collecting and analyzing data and information, and transferring data and knowledge from one site to another. A very large company reported a *knowledge-based configurator* has accelerated the speed of configuring complex teleconferencing equipment from weeks to less than a day. In another example, a company reported AI tools allow it to gather and present data and proposals to customers much more quickly.

Four reports, including three from major defense contractors, were received as to how AI increased competitiveness in ways other than "productivity." For example, one company stated AI allows it to adapt more readily to fluctuations and changes in business. One defense contractor pointed out that AI technology greatly enhanced the capabilities of weapon systems to handle complex, ill-structured, knowledge-intensive problems. Examples of where AI technology has been applied included autonomous (unmanned) navigation of land vehicles, diagnostic adjuncts to complex systems, decision aids, and context-dependent object recognition.

In another example, a defense firm teamed with a university to use neural networks to discriminate between underwater signals. The third defense contractor reported the use of fuzzy logic blackboard technology and intelligent auto-routing techniques to develop complex intelligent software systems that can reason about large amounts of ambiguous, incomplete data that is then used for problem solving. The firm added that this exposure to AI technology has increased its competitiveness in other corporate divisions that now also use intelligent decision systems.

4. AI INDUSTRY STRUCTURE AND PERFORMANCE

Artificial Intelligence has activities in several different formalized industry classifications under the Standard Industrial Classification (SIC) system. However, in each industry AI is only a tiny indistinguishable part for which statistical data is not retrievable. AI vendors and developers are mostly classified with computer software firms. Additional important supporting activities of the technology are found in the "research" and "consulting" industries. Some firms specialize in only one of these areas; others are multi-functional. A select few, such as IBM or AT&T, cover the full range of activities. The principal SIC industries that capture most AI related activities include the following:

Computer Programming Services - SIC 7371

Prepackaged Software - SIC 7372

Computer Integrated Systems Design - SIC 7373

Computer Consultants - SIC 7379

Research, Commercial - SIC 8732

Research, Noncommercial - SIC 8733

Employment in the AI sector, estimated very roughly at 10,000 individuals in 1993, is distributed unevenly across these industries, and several others not shown, such as Computer Training Services (SIC 8243). AI vendors and developers are normally classified in the computer software sector, particularly "prepackaged software" (SIC 7372), in what would be the "core" AI industry if the technology were mature. However, most firms that sell AI systems also provide consultation services, and some also offer training services, from which a significant part of their revenues are commonly derived.

In comparison with the rapidly growing U.S. software sector (the first three SIC designations above), the AI industry represents only about 2 percent of employment. The software sector employed nearly 435,000 people as of June 1993, up 9 percent from the year before (according to the U.S. Bureau of Labor Statistics). Each of the three segments also posted gains - Computer Programming Services, up 10 percent to 182,000; Prepackaged Software, up 8 percent to 142,500; and, Computer Integrated Systems Design, up 8 percent to 110,500. The research and consultative sectors are so broad and varied in scope, aggregate statistical information on them has no relevance to their AI components.

4.1 Industry Description

Artificial Intelligence technology is combining with Object-Oriented Programming,[9] Computer Assisted Software Engineering,[10] and Client/Server[11] technology and emerging as the "***Knowledge Automation Industry***".[12] This industry is an important component of the information age, playing a unique role in the re-engineering of the modern global corporation. Ideally, knowledge automation brings together the widely dispersed "know-how" and experience contained in the "minds" of a business employees and transforms it into

---

[9]Object-oriented programming is a way of structuring programs so that a particular type of data and the parts of a program that process that type of data are combined. Data and the functions that process them are collectively called an "object." Objects are manipulated as a unit; code and data cannot be separated. Many objects can form the parts of large programs.

[10]Computer assisted software engineering is extremely important in today's environment. CASE is "software" that assists software engineers in the writing of new software applications as opposed to building new applications from scratch or without a clear set of objectives.

[11]Originally clients were "dumb" computer terminals wired to mainframe computers (servers). The mainframe did all the processing based on instructions received from a person at the terminal's keyboard. Hundreds, if not thousands, of computer terminals could be wired to the mainframe. Today, desk top computers store and process information on their own, and client-server architectures refer to wiring all types of computer platforms into a network, where each is both a client and server.

[12]"Knowledge" refers to the ability to take a specific action to achieve a goal on given information. Professor of Philosophy Herbert Dreyfus of UC\Berkeley in his book *What Computers Still Can't Do* (1992) describes some of the limits to automating knowledge with the current state of AI technology. A major goal of AI research is to push these limitations back, which underscores the importance of continued long term funding of basic research. Other countries are pursuing the same goal (e.g., see *Future Technology in Japan: Forecast to the Year 2015*).

a strategic asset. Table 3 offers a tabular depiction of the industry.

Table 3. Components of the Knowledge Automation Industry (leverages the value of an organization's in-house knowledge and experience)	
Artificial Intelligence	makes it easier to manage knowledge and experience assets
Object-Oriented Programming	makes it easier to manipulate code
Computer Assisted Software Engineering	makes it easier to manage software development
Client-Server Networks	makes it easier to connect people and computers in an organization and promotes teamwork

A good example of how knowledge can be automated is provided by Ford Motor Company's acclaimed *Computer Aided Parts Estimating System (CAPE)*, a knowledge-based estimator's assistant used by Ford of Europe. CAPE generates, evaluates, and costs automotive parts manufacturing plans. First applied in 1991, CAPE will eventually model and "simulate" electronically every significant production process that contributes to motor vehicle manufacture. The model includes detailed technical information about each machine, process step, and material in the parts pipeline, as well as pricing and cost data. Hundreds of specialized human experts contributed their "knowledge and experience" to the system. CAPE runs on an IBM mainframe and was developed using Inference Corporation's ART tool (case-based reasoning). It operates in an object-oriented modeling environment.

Considering Ford spends several billion dollars on parts each year (a major competitive factor), the payback in cost savings and increased competitiveness is very great. CAPE has resulted in a 50 percent reduction in the time it takes to estimate the cost of parts, and brought a high level of expertise and consistency across the estimating process that was lacking in the past. This has actually reduced the cost of parts, in some cases up to 30 percent, and effected an overall savings on vehicle production costs. In addition, the system is an excellent training tool for new estimators.

The knowledge automation industry, as shown by the employment numbers, is a tiny yet very influential component of the much greater computer software industry. It fills highly specialized "niches" at the top end of the advanced customized software market. In the near term, as AI evolves, and the trend in corporate computer configurations continues becoming more suitable to the use of AI technology, the niche where AI currently finds itself will undoubtedly be enlarged. In the much longer term, AI technology can be expected to become standard practice.

In depicting the knowledge automation industry, a distinction must also be made between the high end and low end of the AI market. The high end or "leading edge" of the industry develops highly sophisticated AI systems, such as CAPE, that often merge the best of several AI technologies (e.g., expert systems, neural networks, natural language, hyper-text, fuzzy logic, etc.) into a single system. These systems are super-engineered "expert systems" that in a very real sense exhibit capabilities near the human expert level. They may be developed by a consortium of specialized firms and universities, and incorporate the latest AI developments coming out of the AI labs. Leading-edge AI systems are typically developed for and often by major corporations, the Department of Defense and other large organizations where the savings and productivity opportunities are the greatest.

The "lower end" of the market is comprised of vendors focused mostly on prepackaged proven AI technologies, the AI tools market. The tools range from simple to complex, are usually inexpensive, and, as with advanced systems, can include several integrated AI techniques. An AI vendor may simply sell or license an empty "shell". He typically also offers development services at the customer's option. About 200 vendors are believed to be active in various niche markets. The niches take various forms. For example, some systems are targeted at specific industries such as aerospace or medicine, while others run on specific computer platforms, or emphasize specific AI technologies or some other special capabilities. A goal of many vendors is to push their product offerings into as broad a market as possible.

The two segments are not as clearly differentiated as they might at first glance appear. For example, many vendors are active in both areas. Also, most vendors keep a close watch on breaking research developments to incorporate new techniques into their products. And some vendors have research facilities of their own, and frequently employ, and in some cases are owned or managed by, university professors or research scientists. In addition, major corporations such as DuPont, General Electric, General Motors, Boeing, Digital Equipment, and others like them have established internal AI divisions, or equity interests in AI vendors, or both. For instance, General Motors reported that the firm had 23 internal AI projects going on in several different divisions of the company. Estimates have it that 70-80 percent of the Fortune 500 firms now use AI technology to varying degrees. However, it is not known how many actually have internal groups specialized in AI, although many do. In many cases the AI people are commingled with other individuals in the corporate management information divisions and are not counted separately.

Another factor clouding the issue is that AI can no longer be considered a "stand-alone" technology, and to an extent is losing identity as a distinct product in the marketplace. The technology is increasingly embedded or integrated in multi-purpose software packages where it can, for example, increase the productivity, performance, and user-friendliness of the package. Also, as AI technology progressed, many of the techniques initially developed as AI technologies lost identity with their AI parentage. These techniques are no longer thought of as AI, but as an integral components of general use software. Instances of this phenomena grow with each passing year. Examples include windowing, spreadsheets, spelling and grammar checkers, E-mail, hyper-text, income tax software, and chess games.

4.1.1 Modeling the Industry - With these caveats, a structural model of the AI industry can be constructed with three closely connected and overlapping components - research, commercialization, and applications. The *research component* is comprised of all entities that conduct basic and/or applied research into AI technology. This includes universities, research companies or "think tanks", and corporate and government labs.

The *commercialization component* packages the technology into useful forms and essentially transfers the technology to end users for a price (or cost if done internally). This component includes AI vendors, internal corporate and government AI groups, consultants, and educational services organizations, including in its broader context the computer science departments at universities that offer courses in AI. Revenues received (or costs incurred in the development of AI) by this component constitute what may by considered the "primary" AI market.

The *applications component* is made up of businesses and government agencies that purchase AI technology as capital assets to automate "knowledge and experience." The applications component mirrors the primary AI market. However, looking beyond that, its true value lies in the increased competitiveness conferred on businesses using the technology. This value is difficult to quantify, but substantial. The published literature on productivity gains resulting from the application of AI technology is voluminous and growing.

For example, Ed Feigenbaum in his book *The Rise of the Expert System Company* (1988) cites numerous examples of productivity gains. He believes the technology is "revolutionary," in that the realized gains can typically be more than a "10-fold" increase over conventional methods. But again, AI is not a stand-alone technology. Its productivity is difficult to measure. Computer hardware technology and configurations and other software technology elements should also be considered as contributing to the gains. More importantly, a good management and work force attitude has an important, if not a determinative, influence on these results.

In a more conventional industry model, these components would form three distinct areas with less overlap than the AI industry currently exhibits; that is, conventional industries are mature in their life cycles. However, AI is still taking shape. The relative size of the components is gradually shifting from research to applications as the life cycle moves up the maturation curve.

To put the applications market in perspective, the business world is changing from one based on mainframe computers to one based on multi-platforms. With the enormous jumps in computing power brought on by advances in chip densities, the cost of computing on PCs and workstations is now much less than mainframes. The proliferation of smaller computer platforms throughout business organizations has scattered ("decentralized") valuable information throughout businesses and other organizations. The phenomena has led to confusion about what hardware and software will be needed, and hesitancy on the part of many corporations to plunge into major investment commitments for information technology. This may take a few years to sort itself out, although some major corporations are moving aggressively to re-invent their organizations.

The challenge of the 1990s and beyond is to extend and integrate knowledge. This includes the integration of information and data management technologies to the point where any level problem can be solved successfully by having access to it in an automated system. As noted, technologies such as AI systems, object oriented programming, client-server architectures, and computer assisted software engineering are key elements in this emerging environment.

Given the description of the industry, can the market for AI technology be estimated? Based on reports submitted to the Commerce Department from universities, firms, and government agencies, estimated AI research expenditures totaled about $206 million in 1993.[13] About half that total was for basic research, performed primarily at universities or private research

[13]The $206 million excludes corporate internally funded research, for which accurate numbers are not available.

centers, mostly under government contract. Estimates of the "primary" or commercial market can be hazardous because of so many unknowns. However, the January 1994 issue of the Harmon Associates' *Intelligent Software Strategies* (ISS) newsletter, estimated the 1993 commercial AI market, which includes both public and private sector sales, at $601 million. This estimate contained $45 million in U.S. Government funded research by "contractors". Since the $45 million is already allocated to our research component, a modified estimate of the ISS commercial market would equal $557 million, or roughly 2.7 times the estimated research total.

4.2 AI Research Funding

Several major AI universities and companies provided the original foundation of AI research, although their role has diminished over time. These institutions performed much of the research and developed AI technology to a state where commercialization was feasible. In many cases, they also provided the professor-entrepreneurs who founded many AI start-up companies in the late 1970s and early 1980s, and initially commercialized the technology.

The original "core" AI research universities included Carnegie-Mellon University (CMU) in Pittsburgh, the Massachusetts Institute of Technology (MIT), and Stanford University. Joining these institutions were companies, sometimes called "think tanks," heavily engaged in AI research. Preeminent among these was the Stanford Research Institute, Inc. (SRI), which has been considered the equal of the core universities almost from the beginning. Other important companies included Bolt, Beranek, and Neuman, Inc. (BBN), AT&T, Xerox's Palo Alto Research Center (PARC), the RAND Corporation, and of course, IBM's scientific centers.

Today, the list of high-powered research organizations and individuals has expanded to many more, as the technology attracted increased attention during the 1970s and 1980s. An additional push came in the 1980s from the Strategic Computing Initiative (1983-1992). At its peak this initiative, which many define as the American response to the Japanese Fifth Generation Project, was spending over $100 million a year, of which about 20 percent (or more depending on how AI is defined) was for AI technologies.[14]

[14] *IEEE Expert*, June 1991, p. 8

The top research institutions handle over $10 million of AI research on an annual basis. Most, but not all, are heavily involved with the Department of Defense. Currently, the research at the top five research organizations represents almost 40 percent of all AI research, and roughly half the total "basic" AI research that takes place in the United States.

Based on survey information collected from government agencies and private academic and commercial sources, estimated funding for AI research from all sources has held steady at somewhat more than $200 million since 1990 (in constant 1993 dollars - per GNP deflator).

> An unknown portion of this $200 million value was actually "development" funding rather than research. Estimates of the development portion based on company/university survey results would place it at about 5-10 percent. Based on ARPA's report, the development portion may by considerably higher. This will be discussed in more detail presently.

The Department of Defense supplied about 60 percent of this total. The Department's Advanced Research Projects Agency provided roughly two-thirds of the Defense portion. Other Federal agencies, primarily the National Science Foundation (NSF) and the National Aeronautics and Space Administration (NASA), provided another 15 percent. Altogether, Federally funded AI research accounted for about 75 percent of total AI research for the period 1989-1994.

The bulk of the remaining 25 percent was funded by U.S. manufacturing and other U.S. businesses, and by foreign sources. Funding by foreign sources varied from as low as 6.5 percent in 1989, to a high of 11.6 percent in 1990. Other funding sources, such as foundations, were not very significant.

Figure 1. Federal Share of AI Research
(in percent)

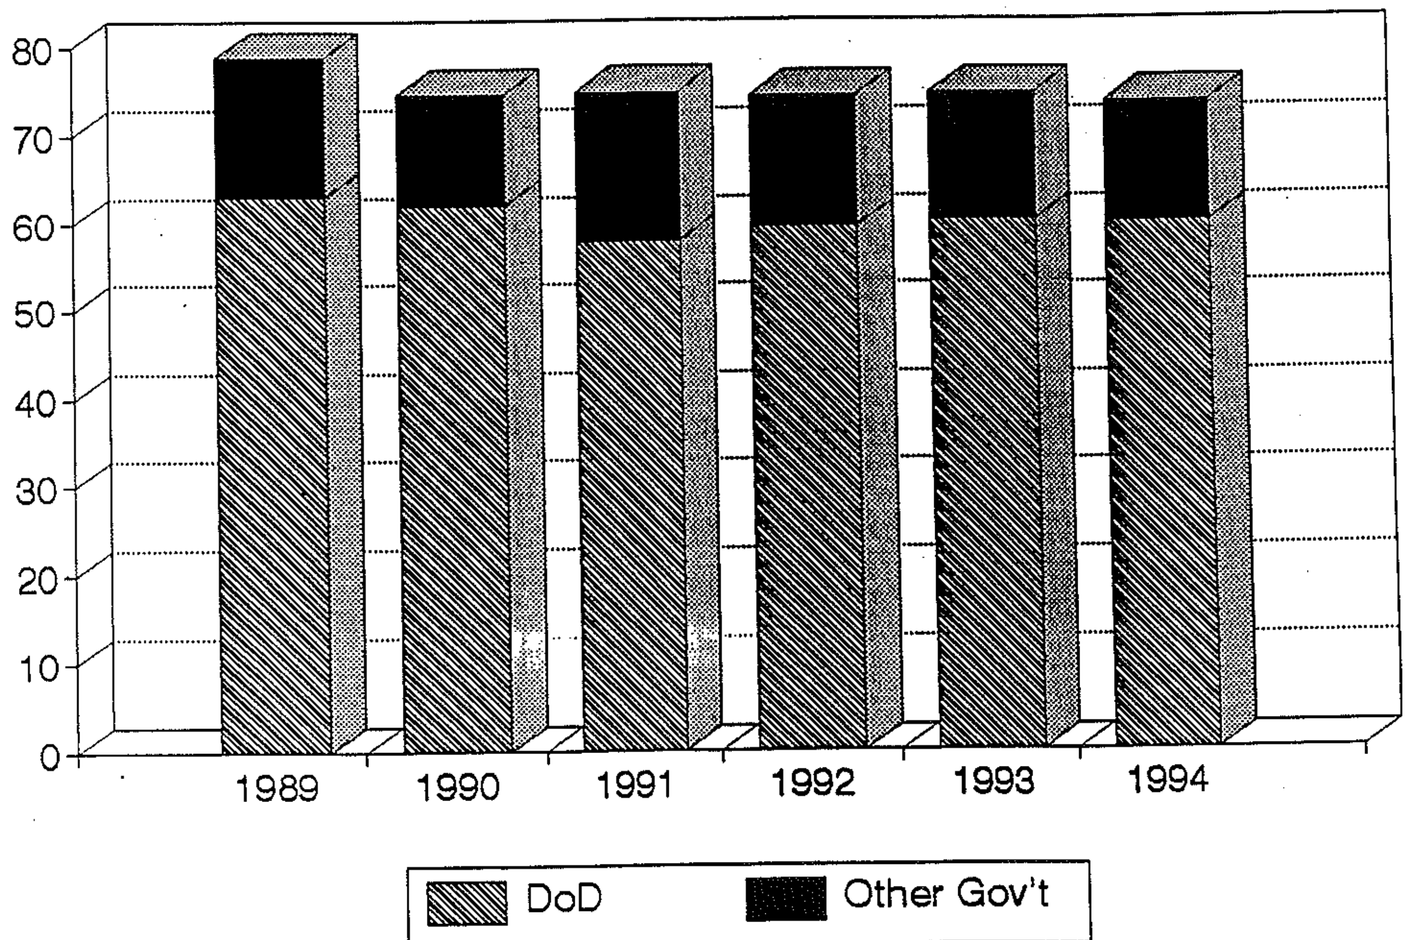

Tables 4, 5, and 6 present estimates of AI research by funding source from 1989-1994. The years 1993 and 1994 were projected by survey respondents.

TABLE 4. ARTIFICIAL INTELLIGENCE - ESTIMATED TOTAL RESEARCH By Funding Source, 1989-1994 (in thousands of constant 1993 dollars)						
FUNDING SOURCE	1989	1990	1991	1992	1993	1994
Dept. of Defense	120,454	131,500	119,124	123,637	124,213	125,345
Other Federal Govt	30,015	27,004	34,969	30,771	29,736	28,894
State and Local Govt	2,260	2,591	2,172	2,949	2,271	2,241
Foreign Govt	1,596	1,846	1,846	2,029	0	0
Manufacturers-US	20,423	22,442	24,604	16,184	16,569	17,120
Manufacturers-Foreign	9,993	14,006	11,296	11,886	8,212	7,615
Other Business-US	4,363	4,030	3,955	12,018	15,306	16,396
Other Business-Foreign	0	6	0	170	111	73
Joint Ventures-US	0	0	80	76	1,258	2,328
Joint Ventures-Foreign	1,025	8,651	7,713	6,820	7,494	8,243
Foundations-US	930	406	265	1,363	1,301	1,183
Total (all categories)	$191,060	$212,483	$206,024	$207,902	$206,471	$209,439
percent distribution (%)						
Dept. of Defense	63.0	61.9	57.8	59.5	60.2	59.8
Other Federal Govt	15.7	12.7	17.0	14.8	14.4	13.8
State and Local Govt	1.2	1.2	1.1	1.4	1.1	1.1
Foreign Govt	0.8	0.9	0.9	1.0	0.0	0.0
Manufacturers-US	10.7	10.6	11.9	7.8	8.0	8.2
Manufacturers-Foreign	5.2	6.6	5.5	5.7	4.0	3.6
Other Business-US	2.3	1.9	1.9	5.8	7.4	7.8
Other Business-Foreign	0.0	0.0	0.0	0.1	0.1	0.0
Joint Ventures-US	0.0	0.0	0.0	0.0	0.6	1.1
Joint Ventures-Foreign	0.5	4.1	3.7	3.3	3.6	3.9
Foundations-US	0.5	0.2	0.1	0.7	0.6	0.6
Total (all categories)	100.0	100.0	100.0	100.0	100.0	100.0

Source: U.S. Dept. of Commerce, BXA/OIRA AI Sector Survey

TABLE 5.
ARTIFICIAL INTELLIGENCE - ESTIMATED TOTAL RESEARCH BY COMPANIES
By Funding Source, 1989-1994
(in thousands of constant 1993 dollars)

FUNDING SOURCE	1989	1990	1991	1992	1993	1994
Dept. of Defense	58,282	70,159	62,356	63,919	61,778	62,825
Other Federal Govt	12,094	9,474	12,698	10,030	10,042	9,526
State and Local Govt	0	0	0	1,128	0	0
Foreign Govt	1,551	1,427	1,588	1,755	0	0
Manufacturers-US	12,243	14,456	16,786	8,844	9,688	10,939
Manufacturers-Foreign	5,892	9,988	7,144	7,773	4,543	4,083
Other Business-US	1,551	1,427	1,360	9,165	12,337	13,636
Other Business-Foreign	0	0	0	75	72	0
Joint Ventures-US	0	0	0	0	1,195	2,268
Joint Ventures-Foreign	0	0	0	0	1,195	2,268
Foundations-US	930	285	265	251	239	227
Total (all categories)	$92,543	$107,216	$102,198	$102,940	$101,090	$105,772
percent distribution (%)						
Dept. of Defense	63.0	65.4	61.0	62.1	61.1	59.4
Other Federal Govt	13.1	8.8	12.4	9.7	9.9	9.0
State and Local Govt	0.0	0.0	0.0	1.1	0.0	0.0
Foreign Govt	1.7	1.3	1.6	1.7	0.0	0.0
Manufacturers-US	13.2	13.5	16.4	8.6	9.6	10.3
Manufacturers-Foreign	6.4	9.3	7.0	7.6	4.5	3.9
Other Business-US	1.7	1.3	1.3	8.9	12.2	12.9
Other Business-Foreign	0.0	0.0	0.0	0.1	0.1	0.0
Joint Ventures-US	0.0	0.0	0.0	0.0	1.2	2.1
Joint Ventures-Foreign	0.0	0.0	0.0	0.0	1.2	2.1
Foundations-US	1.0	0.3	0.3	0.2	0.2	0.2
Total (all categories)	100.0	100.0	100.0	100.0	100.0	100.0

Source: U.S. Dept. of Commerce, BXA/OIRA AI Sector Survey

TABLE 6. ARTIFICIAL INTELLIGENCE - ESTIMATED TOTAL RESEARCH BY UNIVERSITIES By Funding Source, 1989-1994 (in thousands of constant 1993 dollars)						
FUNDING SOURCE	1989	1990	1991	1992	1993	1994
Dept. of Defense	62,172	61,341	56,768	59,718	62,435	62,520
Other Federal Govt	17,921	17,530	22,271	20,741	19,694	19,368
State and Local Govt	2,260	2,591	2,172	1,821	2,271	2,241
Foreign Govt	46	420	259	274	0	0
Manufacturers-US	8,180	7,985	7,817	7,340	6,881	6,181
Manufacturers-Foreign	4,101	4,018	4,152	4,113	3,669	3,532
Other Business-US	2,812	2,603	2,595	2,853	2,969	2,760
Other Business-Foreign	0	6	0	94	40	73
Joint Ventures-US	0	0	80	76	63	60
Joint Ventures-Foreign	1,025	8,651	7,713	6,820	6,298	5,975
Foundations-US	0	121	0	1,112	1,062	956
Total (all categories)	$98,516	$105,266	$103,826	$104,963	$105,382	$103,667
percent distribution (%)						
Dept. of Defense	63.1	58.3	54.7	56.9	59.2	60.3
Other Federal Govt	18.2	16.7	21.4	19.8	18.7	18.7
State and Local Govt	2.3	2.5	2.1	1.7	2.2	2.2
Foreign Govt	0.0	0.4	0.2	0.3	0.0	0.0
Manufacturers-US	8.3	7.6	7.5	7.0	6.5	6.0
Manufacturers-Foreign	4.2	3.8	4.0	3.9	3.5	3.4
Other Business-US	2.9	2.5	2.5	2.7	2.8	2.7
Joint Ventures-US	0.0	0.0	0.0	0.1	0.0	0.1
Joint Ventures-Foreign	0.0	0.0	0.1	0.1	0.1	0.1
Foundations-US	1.0	8.2	7.4	6.5	6.0	5.8
Foundations-Foreign	0.0	0.1	0.0	1.1	1.0	0.9
Total (all categories)	100.0	100.0	100.0	100.0	100.0	100.0

Source: U.S. Dept. of Commerce, BXA/OIRA AI Sector Survey

Note that the share of overall AI research undertaken by universities and firms is about equal. Universities received somewhat more Federal funding; the difference accounted for by the greater amount universities received from non-Defense agencies, particularly the National Science Foundation. In general, universities received between 75-80 percent of their funding from the Federal Government. They also received almost all reported state funding. The companies collectively reported between 70-75 percent of their research funding came from the Federal Government. Consequently, companies received larger relative shares of both commercial and foreign sponsored AI research.

Another source of funding is corporate in-house, or "internally" funded AI research. Although not shown on the above tables, corporate internally funded research is significant, but comprehensive statistics on this research are difficult to gather. Large corporations such as IBM, AT&T, Apple Computer, DEC, Intel, General Electric, Motorola, Ford Motor, and others like them are major players in this area, and each has a substantial in-house capability. The Department of Commerce AI survey provided only a small glimpse of this internal market. The numbers assembled from the survey are obviously grossly understated because so few responses (i.e., from major corporations) were received. An examination of several firms' annual reports, government reports from NSF and the Bureau of the Census, and conversations with experts proved futile; information about in-house research is not available.

Moreover, most of the survey respondents were AI vendors, as opposed to large corporations. While the vendors typically spent over 10 percent of their revenues on research, much of it was actually for "development" rather than research. This could not be segregated and served to overstate reported internally funded "research" values by these vendors. The numbers shown on the table below are the internal research totals "as reported" by 28 companies. (Five were major corporations; their 1993 total was $22.6 million.)

Table 7.
Reported Private Internally Funded Research and Development, 1989-1994
(in thousands of dollars)

Internally Funded Research	1989	1990	1991	1992	1993	1994
then year dollars	$30,147	$27,640	$36,076	$30,329	$33,251	$37,924
constant 1993 dollars (GNP deflator)	$34,270	$30,165	$37,991	$31,062	$33,251	$36,819

Source: U.S. Dept. of Commerce, BXA\OIRA AI Sector Survey

4.2.1 Government Funded Research - Statistical information was collected from Defense agencies - the Advanced Research Projects Agency (ARPA), the Office of Naval Research (ONR), the Air Force Office of Scientific Research (AFOSR) and several Air Force Labs, and the Army Research Office (ARO) - and from NSF and NASA. These agencies provided statistics on the number and value of on-going AI research projects, the starting and ending dates of those projects (when available), and the research organizations that received funding for each project. These agencies, DoD, NSF, and NASA, account for at least 90 percent of the total Federal Government AI research expenditures.

No single organization has assumed greater financial responsibility for artificial intelligence research in the United States than the Department of Defense. ARPA (formerly the Defense Advanced Research Projects Agency), organized in 1958 to fund long term scientific "high-tech" research projects, has financed much of the AI research that has been conducted at American universities, research organizations, and private corporations.

Well known ARPA-sponsored projects included Shakey, SRI's mobile robot, and SUR, the Speech Understanding Research project conducted by several organizations including CMU, BBN, and SRI. ARPA was also responsible for a computer telecommunication network, called ARPANET, that allowed AI researchers to exchange ideas quickly and efficiently. The Strategic Computing Program was announced by ARPA in 1983.[15] ARPA also began major funding of neural networks in 1988. As part of the Department of Defense, ARPA looked for technology with potential military application. Nonetheless, the agency allowed researchers latitude within the boundaries of specific projects.

Most military projects range from 1-3 years; some are for longer periods. In its submittal to the Department of Commerce, ARPA reported 282 on-going projects valued at $209 million. Of these, 132 projects worth $105 million went to universities; 104 valued at $95 million went to private firms; and 46 valued at $8.8 million went to other government agencies. ARPA indicated it plans to spend a total of about $179 million in the AI area during the next three fiscal years (1994-1996).

---

[15]In October, 1983, ARPA issued the Strategic Computing Plan, which contained the following quote: "As a result of a series of advances in artificial intelligence, computer science, and microelectronics, we stand at the threshold of a new generation of computer technology having unprecedented capabilities. The United States stands to profit greatly both in national security and economic strength by its determination and ability to exploit this new technology."

ARPA's $209 million included $27.6 million labeled as basic research, and $70 million labeled as applied research. The remainder, $113.7 million, was not labeled and presumably includes both research and development. The amount for development is not clear. The major research organizations (CMU, MIT, Stanford, and SRI) reportedly received about 35% of ARPA's total funding; 90% of it was unlabeled. Comparing survey responses with that actually reported by ARPA leads to the conclusion that most of the unlabeled portion must actually be "research" or the numbers would not balance.

Although the AI data collected for this study cannot be used to validate a decline in basic AI research, a noticeable shift from basic to applied research has occurred during the last ten years in Defense funding in general and in computer science in particular. Supported by statements by several university officials, this same trend is likely to have also occurred in AI research. Pressures to show results mounted as the military budget was cut back about 40 percent since its peak in 1986 (See Tables 8-13 in Appendix D). Basic research is categorized as high-risk and long-term. It is often the first area cut when budgets tighten.

In a separate report to the Department of Commerce, ONR indicated 45 projects (for $63 million) it managed were funded through ARPA. On its own account, ONR reported 70 projects funded for $36 million. ONR research funding appeared to be dropping off sharply from the late 1980s. For example, 1990 budgeted funding was $9.4 million; in 1991, it dropped to only $1.6 million. In addition, of the ARPA funded research reported by ONR, $17.4 million was budgeted in 1990, and $9.4 million in 1991.

The Air Force reported 101 projects for $71 million. About half the Air Force total was funded through the Rome Lab at Griffen AFB in New York. Most of the remainder comes from Wright Lab and AFOSR, both at Wright-Paterson AFB in Ohio. The Armstrong Lab at Brooks AFB in Texas also funds a small amount that incorporate off-the-shelf AI technology for training purposes. Wright Lab was the only Air Force unit to provide funding by fiscal year. Wright Lab's funding peaked at $5.3 million in 1991, and since has fallen to about $1 million. The Army reported it spent about $2.5 million a year, and intends to maintain that level.

NASA reported 21 projects over the next several years worth about $39 million. All of the NASA funding was reported as going to one of several NASA research centers, particularly the Ames Research Center near San Francisco. Ames, and the others, parcel out some of the money to universities and firms, but the amounts and recipients were not indicated. Nearly all of NASA's projects are for five years or more.

NSF typically provides one-year grants, although some are longer, or may be extended for longer periods. NSF publishes a yearly handbook, which documents each grant in detail. NSF issues over 300 grants a year. Most are small. The average grant is between $50-100 thousand. Some industry observers believe these grants would be more effective if issued in larger amounts for longer terms. About 90 percent of NSF funded AI research is basic in nature. The value of all grants rose steadily from nearly $15 million in 1986, to about $26 million in 1992. Over the three years 1990-1992, NSF issued grants to 163 different organizations for about $63 million.

Table 8.
TOP 15 FEDERALLY FUNDED AI RESEARCH ORGANIZATIONS, 1991

Research Organization	Amount	Percent of Total	Accum. Percent
U\ Southern California	$14,586,504	10.62%	10.62
BBN Lab, Inc.	12,106,797	8.81	19.43
Carnegie Mellon University	9,097,159	6.62	26.05
MIT\MIT Lincoln Lab	9,051,118	6.59	35.64
U\ Massachusetts	5,919,146	4.31	36.95
General Electric	4,234,097	3.08	40.03
Stanford U	3,609,445	2.63	42.66
NASA-Ames Research Center	3,505,000	2.55	45.21
Northwestern U	3,359,419	2.45	47.66
Dragon Systems, Inc.	3,056,312	2.22	49.88
IBM	3,022,179	2.20	52.08
U\ Pennsylvania	3,017,144	2.20	54.28
SRI International	2,929,939	2.13	56.41
Mitre	2,354,443	1.71	58.12
Hughes Aircraft	2,063,215	1.50	59.62
Sub-Total	$81,911,917	59.62	-
Grand Total	$137,370,759	100.00	100.00

Source: Reports received from DoD, NSF, and NASA

Note that SRI is lodged in 13th place on Table 8. This is an anomaly. The firm received greater funding in surrounding years. Overall, SRI received about $25 million of ARPA's $272 million budget which includes $63 million through ONR. This does not include work SRI is doing for the Post Office, NASA and other Federal agencies.

An effort to allocate these projects on an annual basis revealed a peak year of 1991, when funding by these agencies collectively totaled about $137 million. The $137 million was distributed to 195 separate research organizations, including $10.2 million to 17 Federal Labs. In addition to these labs, 116 universities and colleges received $78.7 million; 53 private firms received $47.2 million; and 9 associations (or institutes) received $1.3 million. The major single sponsor was ARPA at over $80 million, or roughly 60 percent of the total. The NSF was a distant second at about $23 million in 1991, or almost 17 percent.

The top 30 research organizations accounted for just over 75 percent of total Federally funded research. The top 15 accounted for almost 60 percent, or about $82 million of the $137 million Federal total. The University of Southern California led the pack with $14.6 million (10.6 percent), followed by BBN ($12.1 million) and Carnegie Mellon ($9.1 million). The top 5 accounted for 37 percent of the total, and the top 10 for 50 percent. Table 8 shows the top 15 organization for 1991.

A major difference between university and commercial research is the much greater focus on "basic" research at universities. This is shown on Figure 2 and Table 9. Universities are usually engaged in the search for new knowledge, and approach problems from a scientific basis. Corporate research is predominantly exploratory in nature, and if found promising, often leads to "development" (or prototyping) of new processes or products. Corporate research is more of an engineer's task, where university research is more research scientist-oriented.

Respondents to the Department of Commerce AI survey bear witness to this difference. In the aggregate, companies reported on average about one-third of their research was "basic," while universities reported more than 76 percent for that purpose. About 40 percent of the company reported research was "internally" funded. Looked at from an individual firm basis, as noted earlier, the great majority of smaller firms actually had little or no internally funded "basic" research. In the case of small AI firms, research was almost totally of an "applied" nature, or used for development expenditures. Several major corporations with substantial in-house AI groups reported nearly all the basic research undertaken "internally" by firms. Company research funded by the Government, however, was the greater share of total company basic research. Percentage results are reflected by Figure 2 and Table 9.

Figure 2. Company Research Focus, 1991

% Basic and Applied AI Research by Area

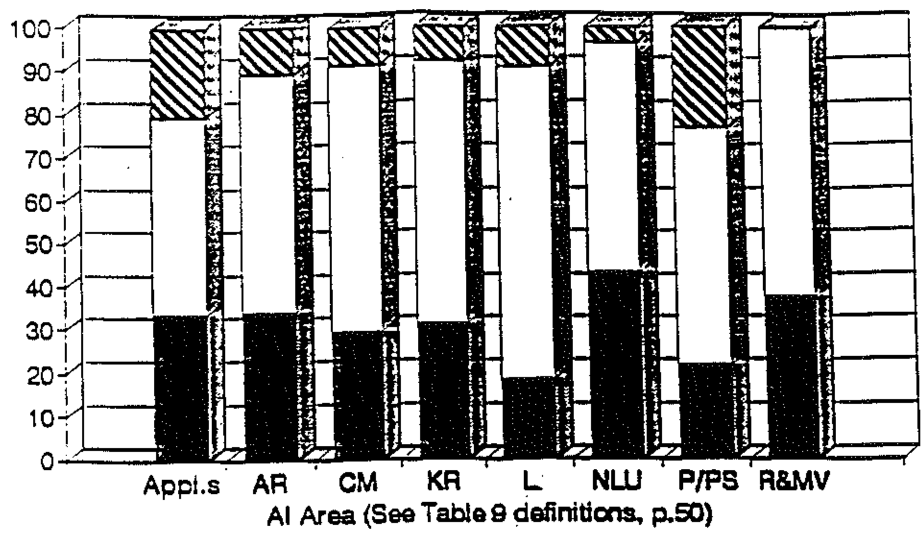

University Research Focus

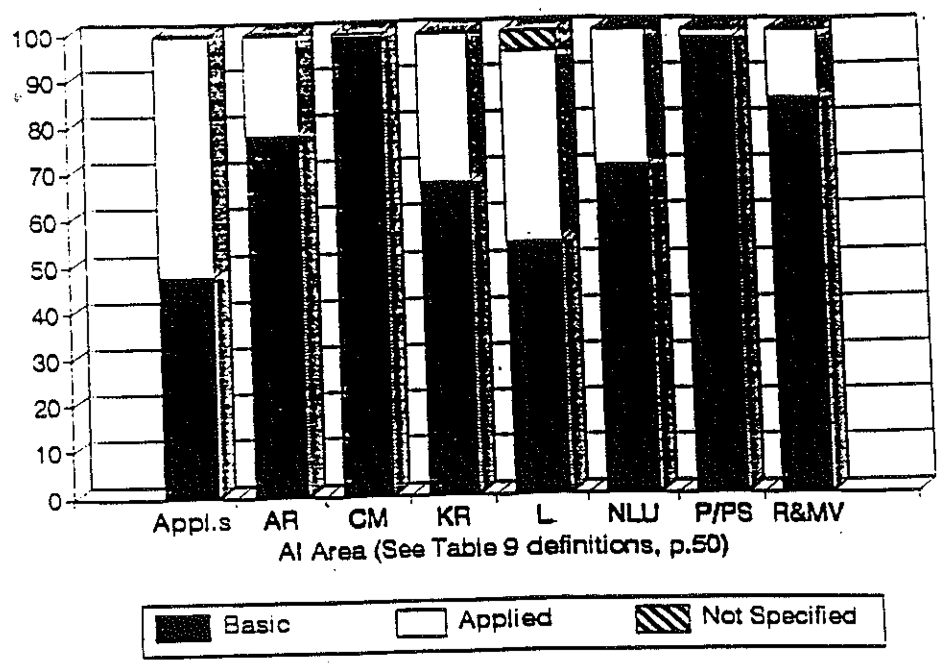

TABLE 9. PERCENT BASIC, APPLIED, AND OTHER AI RESEARCH BY AREA, Projects active in 1991 (comparison of companies and universities)						
RESEARCH AREA	RESEARCH BY COMPANIES			RESEARCH BY UNIVERSITIES		
	BASIC	APPLIED	OTHER	BASIC	APPLIED	OTHER
Applications	33.7	46.0	20.3	47.8	52.2	0.0
Automated Reasoning	34.1	55.2	10.7	78.3	21.7	0.0
Cognitive Modeling	29.6	61.5	8.9	99.4	0.6	0.0
Knowledge Representation	31.3	60.9	7.8	67.6	32.2	0.1
Learning	18.5	72.3	9.2	54.5	41.0	4.5
Natural Language Understanding	43.3	53.0	3.7	70.9	29.1	0.0
Planning and Problem Solving	21.5	55.2	23.3	98.3	1.7	0.0
Robotics and Machine Vision	37.4	62.6	0.0	85.3	14.7	0.0
Total (all categories)	30.8	59.5	9.7	76.3	23.3	0.5

Note: "Other" was not clearly defined. It is assumed to be "development" funding.
Source: U.S. Dept. of Commerce, BXA\OIRA AI Sector Survey

The AI survey also included questions on the employment of research scientists and engineers, from which further comparisons could be made. Table 10 shows the ratio of research scientists to engineers and how it differs between firms and universities. For this exercise "think tanks," such as SRI, were included with the universities since their ratio for this indicator was about the same, and they are not "firms" in the same mold as the others. Think tanks would have been listed separately, except to do so would reveal proprietary information. The firms were divided further into larger (20 or more employees), and smaller (fewer than 20 employees). Also, the data base itself does not cover the entire universe.[16] With these adjustments, the table shows the greater emphasis on research scientists at universities, which simply further confirms the emphasis on basic research at universities. Moreover, no difference in the indicator was found between smaller and larger firms.

Table 10. Ratio of Research Scientists to Engineers at AI Research Activities, 1991						
Type Institution		Research Scientists		Engineers		Ratio of Research Scientists to Engineers
	total reporting	number reporting	number employed	number reporting	number employed	
Larger Firms	14	9	93	11	181	.51
Smaller Firms	26	17	27	22	53	.51
Universities	24	24	266	12	74	3.59

Source: U.S. Dept. of Commerce, BXA/OIRA AI Sector Survey

4.2.2 University/Company Reported Research - Twenty-six universities and 38 companies reported a total of 408 on-going (circa 1991-1992) research projects valued at $347 million in response to the survey. The firms reported 218 projects valued at $203 million, and the universities, 190 projects worth $144 million. Figure 3 and Table 11 below present the allocation of research by AI area. Individual projects sometimes included more than one area of research. These tended to be larger than the single area projects. In cases where this occurred, total dollar amounts were proportioned equally to each area referenced. Of the total 408 projects, 248 (60 percent) involved just one area of research. Another 61 projects involved just 2 areas; and 99 involved 3 (or more) areas.

[16]An additional caveat deals with how respondents interpreted the question. Various inconsistencies appeared in the way the question was answered. For example, one entity reported a number of knowledge engineers, but left the "engineers" block blank. Another reported everybody employed, including "trainees" and clerical individuals, as research scientists. Other organizations, while reporting substantial research, left both blocks blank. These anomalies were removed from the analysis.

Figure 3. AI Research by Area

Comparison of Companies & Universities

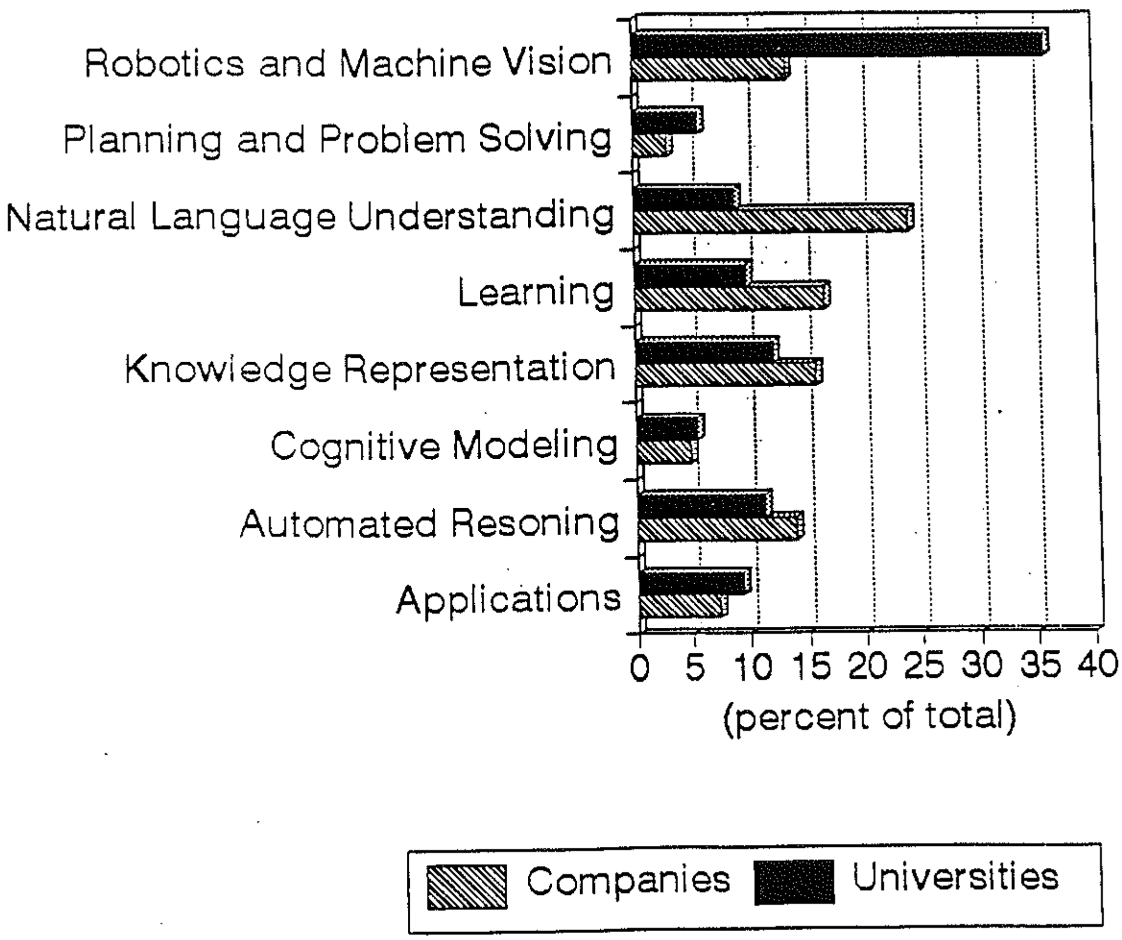

TABLE 11. PERCENT AI RESEARCH BY AREA OF RESEARCH (comparison of companies and universities)		
Research Area	Companies	Universities
Applications	7.4%	9.5%
Automated Reasoning	14.0	11.4
Cognitive Modeling	5.0	5.5
Knowledge Representation	15.9	12.2
Learning	16.7	10.0
Natural Language Understanding	24.2	9.0
Planning and Problem Solving	3.1	6.0
Robotics and Machine Vision	13.7	36.4
Total (all categories)	100.0%	100.0%

Source: U.S. Dept. of Commerce, BXA\OIRA AI Sector Survey

Universities show a heavy concentration in robotics, which is one of the high focus areas of Defense and the major AI universities. Firms are focused more on natural language understanding, particularly "speech recognition," which has made major strides in recent years. For example, a vendor stated that prices for speech systems fell 10-fold while machine vocabulary increased from a few thousand to over 30,000 words between 1990-1994.[17] Natural language, which is making computers easier to use (and thus more valuable), was also pushed by Defense.

[17]Conversation with Dragon Systems representative at Trade Fair in March, 1994.

4.2.3 Employment at Universities - Table 12 contains various employment measures at reporting AI universities.

Table 12. Employment Profile of Universities					
Occupation Title	number of responses	1989	1990	1991	1992
Knowledge Engineers	8	15	22	37	40
Software Developers	11	122	137	156	157
Research Scientists	23	209	218	240	240
Graduate Students	11	286	317	304	362
Clerical/Admin.	12	59	68	77	81
Total	24	691	762	804	880
# Phd's in above Total	21	161	172	194	203
# of Trainees	17	267	294	279	316
Reported Training Expenses ($000s)	14	$3,055	$3,251	$3,184	$3,532
Reported Payroll ($000s)	20	$16,813	$19,615	$22,933	$23,520
# of Computer Scientists	24	573	612	613	679
# of Mathematicians	1	2	2	2	1
# of Cognitive Scientists	5	15	25	29	27
# of Engineers	10	47	52	64	63

Source: U.S. Dept. of Commerce, BXA/OIRA AI Sector Survey

AI-related employment at 22 universities rose between 1989-1992, from 691 to 880, up about 27 percent. Nearly all of the increase was accounted for by Northwestern University, Carnegie Mellon, MIT Lincoln Lab, and the University of Washington. Northwestern alone reported more than a 3-fold increase, and accounted for well over half of the numerical increase. Northwestern's Institute for the Learning Sciences, established in 1989, has grown rapidly with funding from Arthur Andersen, IBM, and others. Three schools, with smaller AI programs, also reported declines in employment ranging from 20-35 percent.

The number of research scientists rose 15 percent during the period while graduate trainees and research assistants rose 26 percent. Four institutions reported a drop in the number of research scientists, and two reported a drop in the number of trainees. The rapid rise in knowledge engineers is attributable mostly to one school. The increase in software developers of 29 percent was split between three schools. The number of PhDs rose evenly with total employment, while computer scientists trailed slightly, finishing about 19 percent higher. Expenditures per trainee remained at about $11,000 per annum throughout the period. Payroll rose 40 percent, indication a moderate rise in income.

4.3 AI Commercialization

Product life cycles have six stages as depicted below:

Table 13. Product Life Cycle
Basic Research
Invention
Commercialization
Rapid Growth
Maturity
Obsolescence

AI technology is no different. It too has a life cycle that if graphed, time on x-axis and revenues on y-axis, would be shaped like a stretched-out "S" leaning to the right. Although revenues may begin trickling in during the commercialization phase, costs begin accumulating from the start. The break even point generally will not occur until during the rapid growth phase, after which time profits can be realized. Various AI technologies are represented in all the life cycle phases.

AI is a brain-intensive technology (as opposed to capital intensive) with low start-up costs. Thus, literally anyone that understands the technology can start an AI business. However, the path to success is treacherous. Of the many dozens of AI vendors that entered the business in the late-1970s and early 1980s, only a handful still remain. Also, of those that do remain, most have no resemblance to their original form, and have a completely new set of people in management positions. However, easy entry has allowed a new crop of vendors to appear as the older ones depart. Perhaps not surprisingly, the new vendors are frequently established by individuals bailing out or discharged from failing older firms.

Knowledge-based systems (KBS) were the first AI technology to breakout of the research labs and eventually reach maturity. The first KBS vendors appeared on the scene about 1980, and by 1990, the technology had attained maturity. Natural language programs struggled through the 1980s, and only in the last 2-3 years appear to have entered the rapid growth phase. Neural networks (NN) also struggled for many years, and only started taking off in the late 1980s. By some projections, NNs are expected to grow rapidly (from a small starting point) for the remainder of this decade, and by then exceed sales of KBS systems. Fuzzy logic has not been that popular in the United States, although recently more attention has focused in that direction. Ironically, fuzzy logic is already a mature AI technology in Japan.

Robotics and machine learning are two areas where real progress in research has been slow. The military, well known for developing cutting-edge AI systems, has done a lot to push robotic technologies, and has fielded some systems using "vision" techniques. Rodney Brooks at MIT has taken a bottom-up approach (primitive to advanced) and has achieved results. He has now successfully created the equivalent of artificial "bugs." Genuine machine learning, that is, the kind of learning people do, will require several major breakthroughs, and perhaps new approaches, before it comes out of the research labs.

4.3.1 Commercial Environment - AI commercialization began in the mid-1970s, and is still going on in various areas. When AI technology was first being introduced, it was greeted with much enthusiasm and media "hype" - not uncommon in the American experience. However, this overstatement of the underlying reality led to disappointment when the technology failed to deliver on expectations. But the high expectations and overselling of the technology has not impeded the steady growth in the AI market, which now stands at close to $600 million.

AI experienced market acceptance problems for various reasons - excessive promises, difficult technology, management culture, fear, etc. However, what few realized was that AI technology came out of the laboratory ahead of its time; corporate computing was not ready. Corporate computing in 1980 was mainframe-based, centralized, focused on accounting and payroll, and unwilling to change. Early AI vendors tried to get corporate computing management to re-engineer their computing systems, before corporate managers understood the changes coming about. Since then advances in computer technology have forced change. Now corporate computing is becoming multi-platformed, decentralized, and involves every facet of the business; and it is in a state of change, which is favorable to the use of AI techniques. The future of AI is very optimistic.

The commercialization of AI can be separated into two overlapping phases - introductory and restructuring. The high visibility first phase, or introductory phase (1975-1990), saw a large crop of start-up companies emerge under strong influence from universities and corporate AI research groups. These companies, the AI vendors, focused primarily on marketing knowledge-based systems - mostly as stand-alone expert systems. By the mid-1980s, several dozen AI vendors were formed. This phase also saw the rise and fall of the dedicated machine companies, again with close ties to key research institutions.

In addition, the first phase ushered in a period of intense corporate interest in AI which peaked in the mid-1980s. This may have been a reaction to the Fifth Generation Project in Japan, which sought nothing less than to dominate computer technology in the 1990s (it failed), or a reaction to the advent of powerful small computers then finding their way into corporate America. It could also have been activated by the recent formation of many AI vendors to sell AI products into commercial markets. It was probably a combination of these things and more.

In 1985, Business Week[18] magazine ran an article reporting on this explosion of corporate interest in AI. That year, the article stated, "an estimated 150 companies, including General Electric, Gould, Shell Oil, and 3M, spent more than $1 billion to maintain in-house AI groups." Continuing, the article reported, "The interest is spreading far beyond the leading-edge companies." But the enthusiasm, however great it may have been, soon faded

[18]"Business Week," July 1, 1985, p. 78

as companies realized it was premature.

In developing AI during the introductory phase, some corporations took what may be called the "elephant approach," while others took the "rabbit approach," depending largely on corporate "culture."[19] Corporations taking the elephant approach spent hundreds of thousands or even millions of dollars assembling or hiring a special team to develop a major AI system. These were used, perhaps to solve a previously untractable problem, or, in the case of Schlumberger, for example, to locate promising oil yielding geologic formations.[20]

The elephant approach was often taken by corporations with a centralized hierarchical management organization. Elephants have 22 month gestation periods, which reflects both the scale and time it took to develop these systems, many of which were extremely good. A criticism of the elephant approach is that it involves relatively few people in a centrally controlled effort, and therefore does not transfer the technology to the work force. It is consistent with the management "Theory X" (i.e., euphemistically, "use the whip" to motivate employees), which minimizes creative inputs from staff. In other words, the work force did not learn the value of AI technology.

Corporations taking the rabbit approach (30 day gestation period) take a relatively simple and low cost AI "shell," train many people in the company how to use it, and then encourages them to find applications in their work areas. This was done at DuPont, for example, largely at the instigation of a gentleman named Ed Mahler. Today, DuPont has hundreds of systems deployed that save the company tens of millions of dollars a year.

The rabbit approach is more consistent with "Theory Z" (i.e., people are creative and motivated, just give them the tools and tell them what you want done) and serves to transfer the technology throughout the corporation. A dozen or so AI firms sold or licensed rabbit type systems. For example, in 1987, Paperback Software announced sale of its 10,000th *VP-Expert*. Other firms with easy to use AI systems included Information Builders, 1st-Class, and Borland with its *Turbo-Prolog*.

---

[19]Professor Bob Smith in March 1994 telephone conversation with Commerce Department official.

[20]Schlumberger's *Dipmeter Advisor* reportedly took 11 years from prototype to deployment.

Paul Harmon, a respected industry reporter, wrote that corporations spent a lot of time and effort experimenting with AI and expert systems techniques in the mid-1980s, but have not moved from experimentation to widespread use. There were too many problems with the AI and expert systems techniques offered to corporations at that time. Keep in mind that most companies reviewing AI and expert systems solutions were "mainframe-based" and had no plans to change. AI and expert systems required more power than was available on the hardware in most companies then. In addition, expert systems products were often hard to integrate with the procedural languages that dominated commercial computing. When expert systems vendors talked about things such as client-server networks, the ideas sounded remote from commercial computing as it was practiced in 1985.[21]

4.3.2 AI Vendors - Most of the new AI vendors were formed by, or with the help of AI researchers and professors, and often bankrolled by one or more major corporations (sometimes referred to as "angels").[22] The university people were in a strategic position to recognize commercial opportunities, and some felt it their responsibility and right as developers of various AI technologies to cash in on the results.

For example, Larry Harris founded AI Corporation (AICorp) in 1975. Harris conducted extensive research on natural language processing as a professor in the Computer Science Department of Dartmouth College. His primary product at AICorp was *Intellect*, a natural language program designed in 1979 to "interface" between computers and computer users that allowed communication in English. Intellect was one of the first natural language commercial products.

Also in 1975, Brown University professors Leon Cooper and Charles Elbaum formed Nestor, Inc., the first neural network company. In 1979, Roger Schank, then a Yale professor and

---

[21]*Intelligent Software System Development* (1993) by Paul Harmon and Curtis Hall, p. 22.

[22]The term "angel" is often applied in Japan; it refers to corporate sponsorship of new ventures. However, other than corporate sponsorship, Japan's venture capital market is very restrictive - the big companies call the shots. In the United States the venture capital markets are wide open. Money is raised from a wide range of competing sources, including corporations. As a result, in America more new companies are created each year than in the rest of the world combined. However, most fail within two years. An angel (by definition) will finance the new venture through the difficult early years in anticipation of a future payoff.

author of *The Cognitive Computer*, a book about the fundamentals of AI and natural language processing, founded Cognitive Systems. Schank, now at Northwestern and Director of the Institute for the Learning Sciences, has long promoted the concept of the "entrepreneurial university" in which professors would be encouraged to both remain on the faculty and develop outside commercial interests.

IntelliCorp, starting as an expert systems company, was founded in 1980, by four Stanford scientists, including Ed Feigenbaum. The company was originally named IntelliGenetics, and set up to market biotechnical software. In 1983, the company went public, raising about $9 million. Through its highly regarded *Knowledge Engineering Environment* (KEE) system first introduced in 1983, IntelliCorp created a market for high-end commercial expert system tools.

In 1986, Amoco Oil Company purchased a large equity stake (60 percent) in IntelliGenetics which was subsequently spun off from IntelliCorp. IntelliCorp became a mainframe expert system development company, emphasizing high-level systems to link then current technologies.

In 1987, IntelliCorp and IBM agreed to develop a mainframe version of KEE, whereby IBM would market KEE under the IBM logo. A year later IntelliCorp announced its mainframe version of KEE, called IBM KEE, which resulted from the two companies' agreement. Also in 1987, the company introduced the KEE Connection aimed at customers who needed expert systems capable of gathering information from existing corporate databases. In addition, the company introduced a version of its KEE product that ran on advanced desktop computers based on the Intel "80386" processor.

In 1981, Teknowledge was formed by a dozen or so AI researchers from Stanford, MIT, and other institutions. Advertising itself as "the single source for knowledge engineering technology," Teknowledge was the first company to develop and market expert systems as a commercial enterprise. In 1984, General Motors and FMC Corporation each acquired an 11 percent equity in Teknowledge for $4 million. Soon afterwards, NYNEX and Procter & Gamble purchased equity interests. About this time, Teknowledge developed important AI systems for Northrop to help schedule parts flow in its Hawthorne, California aircraft factory that raised productivity spectacularly.

In 1986, the company stopped licensing LISP (LISt Processing Language) versions of its expert system shell, moving to "C" language - the first company to move away from LISP.

This caused a good deal of consternation in the AI community. After a good year in 1986, sales plummeted, and many executives deserted the company, and the corporate sponsors lost money. In 1989, American Cimflex and Teknowledge merged to form Cimflex-Teknowledge. This proved to be a disaster as sales dropped sharply over the next several years. In 1990, the company stopped selling expert systems tools as separate products, choosing to focus on the AI consulting business and using their *M.1* and *S.1* tools in conjunction with consultation efforts.

Gold Hill Company was founded in 1982 by people from MIT and Wang Corporation to develop LISP languages for IBM PCs. In 1986, the company announced ACORN, an expert-system building environment for the IBM PC. ACORN was designed for experts who were not experienced programmers. The company's main products were Golden Common LISP and GoldWorks. Its Common LISP was one of the first and remains one of the most powerful and diversified development environments for expert systems and other AI products.

The giant accounting firm Coopers and Lybrand used Gold Hill's Common LISP to field one of most famous expert systems called *ExperTAX*. ExperTAX is in daily use by hundreds of Coopers & Lybrand tax consultants, who run the program on portable personal computers. GoldWorks, an expert system shell written in LISP for the IBM PC, was introduced in 1987 at a price of $7,500. A Macintosh version was introduced the next year.

Inference Corporation was also one of the pioneers in commercial expert systems. Inference was founded in 1983, with funding from Control Data. The American Express *Authorizer's Assistant* expert system, developed using Inference's Automated Reasoning Technology (ART) product, is one of the best-known and most frequently cited expert systems. In 1984, Lockheed Corporation purchased a 12.5 percent interest in Inference for $4 million. In 1985, Ford Motor Co. acquired an equity in both Inference and Carnegie Group, Inc. for $28 million. In 1988, Inference introduced a C-based version of its ART expert-system tool called ART-IM (for Automated Reasoning Tool-Information Management). Today, Inference is the leader in case-based reasoning systems.

Carnegie Group was also founded in 1983 in Pittsburgh, drawing on the extensive AI talent available at CMU - and in fact, was partially funded by CMU. However, most of its initial $4 million funding came from Digital Equipment, Boeing, and France's Generale de Service Informatique. These three investors pledged a total of $14 million worth of development contracts over the next four years.[23] Four CMU professors were in from the start, including Raj Reddy, the Director of CMU's Robotics Institute. Carnegie Group initially specialized in the application of AI to manufacturing. The Group also offered a wide range of courses and seminars covering various aspects of AI.

In 1989, Carnegie Group delivered a system to diagnose production process problems with fuel injectors to Ford Motor. In 1990, Carnegie Group became the pivotal member of the *Initiative for Managing Knowledge Assets* (IMKA) consortium, with Ford, Texas Instruments, US West and Digital Equipment Corporation. IMKA developed the Representation of Corporate Knowledge (ROCK) model which provides the functionality for creating applications in a client/server architecture (with object oriented capabilities).

For Ford, Carnegie created the Service Bay Diagnostic System (SBDS). SBDS addressed problems involved with dealer servicing of cars under warranty. The increasing complexity of auto subassemblies had made it easier for mechanics to replace an entire subassembly rather than replace the part causing the problem. This was very expensive for Ford since all of the parts in the subassembly were passed back to the manufacturer.

With SBDS on a PC in every service bay to guide mechanics through the diagnosis and repair procedure, time and money were saved. The mechanics stayed current, and customer satisfaction was increased. Carnegie created an intricate natural language front-end to SBDS to make the system more accessible to Ford's mechanics.[24]

Neuron Data (founded 1984 in Palo Alto, California) brought an object-oriented flavor to the AI market. Two Frenchmen that marketed Apple computers in Europe established Neuron Data. Its primary product *Nexpert-Object* permits development and delivery of sophisticated expert systems across many computer platforms, such as Apple Macintosh, IBM PC and compatible, DEC VAX, and IBM mainframes. This was the first AI system to run on a Macintosh, and it was written in "C."

---

[23] *Business Week*, July 1, 1985, p. 78

[24] *The Brain Makers: Genius, Ego, and Greed in the Quest for Machines that Think.* H.P. Newquist, 1994, p. 441.

Neuron Data went on to design software tools that could be invisibly embedded in other applications. In addition, Neuron offered these tools on almost any platform, from Mac's to mainframes. Companies like Microsoft and the major database vendors started offering AI capabilities in their products. Eventually, all computer programs will have this intelligent capability.[25]

Many other AI companies were formed. To mention a few, in 1979, Advanced Decision Systems was founded. In 1981, Speech Systems and Thinking Machines (by MIT researchers) were started. In 1982, Dragon Systems was founded by IBM voice recognition researchers. Kurzweil, a competitor of Dragon Systems in voice recognition systems, was also established in 1982, by Ray Kurzweil, a former top student of Marvin Minsky at MIT. In 1983, Syntelligence was founded by SRI and Schlumberger research scientists. In 1984, Arity was established by exiting Lotus employees, with Lotus funding, to develop Prolog applications for the PC. In 1986, Artificial Intelligent Technologies was established.

AION was founded in 1984. In 1986, AION introduced versions of its Application Execution and Development Systems (AES and ADS) for IBM mainframes operating in the VM operating environment. This was the opening salvo in the war of the mainframe expert systems, the hottest one going in AI at that time. AICorp also targeted the "mainframe" end of the market, joining the fray. In 1988, the company introduced its *Knowledge-Based Management System (KBMS).*

4.3.3 The Brief Life of AI Computers - Another aspect of the early commercialization period was the establishment of several specialized computer manufacturers offering dedicated AI computers. These were the LISP workstations, which were particularly adept at running LISP programs. Two notable companies in this area, each founded in 1980, were LISP Machine Inc. (LMI) and Symbolics. LMI was founded by Richard Greenblatt and Alex Jacobson. The LMI computer, called Lambda, incorporated an MIT created architecture called NUBUS, and proved very useful for running LISP programs. Symbolics was founded by alumni from MIT's AI Lab. The firm offered the "3600," also a LISP machine and also based on MIT designs, that competed with the less expensive Lambda. Other machines were later introduced by Texas Instruments, Sun Microsystems, and Apple. In 1986, Nippon Telegraph and Telephone entered the market with a LISP machine called ELIS. The new machine featured a Japanese language interface, a Motorola "68010" microprocessor, and a

[25]Ibid., p. 442

speed said to be several times faster than Symbolics 3600 machines.

Although Symbolics reported losses of almost $20 million in 1984, the dedicated machine market skyrocketed. Sales of these specialized computers were estimated to have soared to $364 million in 1985, up 59 percent over 1984. The two top vendors, Symbolics Inc. and LISP Machine Inc., projected their revenues would grow 50-100 percent that year.[26] However, this was very short lived and was destined to collapse as less expensive, and rapidly improving general purpose computers, such as the Unix machine and IBM PCs, encroached on the market.

As expected, the companies ran into financial difficulties caused by the erosion of the market. As this hemorrhaging continued, a number of LMI executives resigned to form Gensym Corporation (1986) and develop expert systems for the financial services markets. In 1987, LMI filed for bankruptcy, while Symbolics reportedly laid off about 160 people, despite revenues exceeding $100 million. Gigamos, a Canadian AI company and distributor of LISP machines, purchased LMI that year to protect its interests. In 1988, Symbolics laid off another 225 people and eventually dropped the line altogether.

4.3.4 AI Vendor Problems - Not everything was rosy among AI vendors either. For example, in 1985, Verbex was shut down by parent company Exxon, which had failed in attempts to sell off the voice recognition unit. (Some years later Verbex was independently resurrected.) In 1986, Breit International founder Bernadette Reiter was removed from the company and Breit was shut down. Originally a very secretive AI offshoot of Martin Marietta and Hewlett-Packard, Breit opted for a more commercial product strategy. Also in 1986, Excalibur, along with its Savvy Retriever, a natural-language query system option for its Savvy PC system, went into bankruptcy proceedings, and shut down about a year later. (In 1991, with funding from Japanese investors, the company was reestablished and is now a serious player in the neural-network market.)

The industry sometimes refers to this period as the "AI Winter," particularly the year 1987. At this time, the fledgling industry saw sales plummet, a reduction of corporate interest, and a shuffling of AI vendor executives between companies. Frey Associates shut down, and most of its executives (including founder Eric Frey) went to MicroProducts, which also purchased some of Frey's technology assets. Silogic filed for Chapter 11 bankruptcy and

---

[26]Ibid.

then shut down. Human Edge filed for Chapter 7 bankruptcy and liquidated. Many of its assets, primarily PC and Macintosh AI tools, were picked up by Human Intellect. Scott Instruments received last-minute funding to avoid bankruptcy. Palladian removed its founder from chairman and CEO positions.

Among the pioneer firms, Inference suffered losses and laid off approximately 20 employees. IntelliCorp laid off 10 percent of its people; reported losses for the second half of the year; and announced a corporate "realignment," which also involved new product strategies. The firm also hired a new president and CEO, Katharine Branscomb, a former executive with AION. AICorp reported more than $3 million in losses for the year ending March, 1987. Teknowledge also lost $3 million. Also, several top executives at Carnegie Group quit, and with former Inference employees formed Intelligent Technology Group to develop intelligent applications for the financial and text-based commercial markets.

In contrast to 1987, because 1988 was the largest single year of AI corporate revenues, reaching about $1 billion. This included all the combined sales of LISP machines (which by far was the biggest single item), expert system tools, natural language products, neural networks, voice recognition systems, consulting and development services, programming languages, and peripheral business areas such as additional hardware needed to support individual systems.[27] Thus, 1988 marked the end of the introductory phase of AI commercialization, where AI was marketed more or less as a stand-alone technology.

4.4 Restructuring, Mergers, and Acquisitions

After 1988, LISP machine sales dropped off rapidly, as did development of AI systems in the LISP language. This was a major pivot point in the market, brought on by significant technical advances in general purpose computers, now capable of running AI programs. The introductory phase was practically over, and a period of restructuring was being ushered in. Most of the larger firms developed AI systems in other languages, such as C, C++, and Common LISP, and adapted them to run on many different platforms.

---

[27] *The Brain Makers: Genius, Ego, and Greed in the Quest for Machines that Think*, by HP Newquist, 1994, p. 438

Firms were targeting different markets, and began to disassociate themselves from the AI label, preferring instead to be called "intelligent software" or "business process automation" vendors. Sales plummeted for many of the vendors, as they reoriented their strategies to the rapid changes taking place in the corporate computing markets. Many more went out of business or were absorbed into the operations of other firms.

In OIRA's survey, AI firms were asked to describe any mergers or acquisitions their organizations were involved in recent years, and what impact the mergers or acquisitions had on their operations. Ten responses were received. The most significant merger to take place during the period occurred on September 30, 1992, between AICorp, headquartered in Waltham, Massachusetts, and AION Corp., headquartered in Palo Alto, California - at the time the number one and two vendors in total sales in the field. The combined operations of the two firms came to about $40 million in 1992, giving the new company a major share (nearly 30 percent) of the vendor supplied knowledge-based systems market. The new company adopted the new name of Trinzic Corporation in part to distance itself from the tainted reputation AI vendors projected in the marketplace and to change its strategic direction. Trinzic went on to acquire Channel Computing in March 1993. The purchase of Channel positioned Trinzic in the general purpose client-server tools market.

The 10 survey respondents' reaction to the AICorp/AION merger was mostly negative, as might be expected from competitors. However, over 40 non-respondents to this question were "silent" on the issue. One firm saw this merger as an indication that the AI market, as defined in the mid-1980s, was now maturing. The vendor stated the merger convinced them that to be successful in this highly competitive market, they needed to view knowledge processing as "only" one of many tools to improve the productivity of software engineers, and as such it will need to be integrated with other software development tools. This firm reported it altered its strategy accordingly.

A second firm indicated they were dislodged from an AION training center because of the merger, leading to a loss of revenues. Another said they were affected by both the reduced number of competitors, and potential customers' increasingly negative perception of the industry. Another firm said the merger decreased the number of tools being developed and made available as commercial-off-the-shelf products. A related general comment made by still another respondent was that mergers affect an acquiring company's capital availability for new investment in a negative way (i.e., capital used for debt financing instead of capital equipment).[28]

IntelliCorp purchased MetaKnowledge in early 1990, a small developer of an object oriented tool called *KAPPA*. The plan was to build market awareness of IntelliCorp as an object-oriented company by focusing on KAPPA and de-emphasizing its well-known expert system development environment - KEE. In late 1991, KnowledgeWare, a leading vendor in computer assisted software engineering (CASE) and a firm with sales exceeding $100 million, was on the verge of purchasing IntelliCorp. However, KnowledgeWare announced a huge quarterly loss, and backed out of the deal.[29] Already in bad shape, this sent IntelliCorp into a tailspin. For the year ended June 30, 1991, IntelliCorp lost almost $12 million on revenues of only about $14 million. The next year saw sales drop to $9.2 million, with losses of $9 million. J.C. Martin, another CASE vendor, acquired an equity interest in IntelliCorp, and although it has not been a smooth ride, IntelliCorp today appears to be on the rebound.

In other actions, Inference Corporation acquired Expertech, Ltd., a U.K. based expert systems developer and distributor in November, 1990. Coral Software was purchased by Apple Computer. In so doing Apple acquired Mac Common LISP, an object-oriented language now commonly used within the Apple AI community. Several years ago, Quintus Corporation took over the marketing of its main competitor for Prolog on the Mac. In November, 1992, General Dynamics sold its electronics division, now totally owned by the Carlyle Group. The new company is named GDE Systems, Inc. Also, Computer Recognition Systems Company acquired Octek in April 1988. Octek was the machine vision division of the Foxboro Co. Booz-Allen acquired ADS in 1991. ADS, located in Mountain View, California, had 165 employees, and 1990 sales were $19.5 million. ADS clients include ARPA, military laboratories, the intelligence community and commercial companies

---

[28]The ill effects of debt financing (a merger) would normally be short lived for the acquiring company. The new merged company's overall stock value may rise if the merger is perceived as beneficial, and even exceed the value of the previous two firms, and thereby improve the capital position. Also, the money is not lost from the economy. The merger probably benefitted the stock holders of the company being acquired, who will reinvest the money gained in other companies, if not back into the new merged company.

[29]*The Brain Makers: Genius, Ego, and Greed in the Quest for Machines that Think.* H.P. Newquist, 1994, p. 422.

in aerospace, transportation and manufacturing.[30]

4.5 Recent Industry Performance

According to an industry newsletter, *Intelligent Software Strategies*, the overall vendor market for knowledge-based system (KBS) development tools more than doubled between 1988-1993, rising from $67.5 million to $143.5 million. This kind of growth is quite remarkable considering the turmoil, bad press, recession, and major restructuring going on in the sector during these years. But the AI systems were getting easier to use; the vendors were gaining experience and getting smarter; and the technology was by this time "proven." Also, the fact that prior to about 1988 customers were often forced to purchase million dollar LISP machines to use AI systems probably inhibited growth.

In addition to KBS tools, several other kinds of AI development tools were also covered by the newsletter in 1993 for the first time. These included neural network tools ($25.9 million), fuzzy logic ($6.2 million), natural language development tools ($25.6 million), and case-based reasoning tools ($6.1 million). Adding these additional categories to KBS, the newsletter estimates total worldwide sales of all intelligent software building tools by North American tool vendors amounted to $207 million in 1993. Most American vendors sell from 20-30 percent of their products abroad.

Earlier reference was made to an even larger overall market for AI systems, which the newsletter estimated at $601 million in 1993.[31] As a very rough estimate, the newsletter also reported the United States made up about 65 percent of the global AI market.[32]

---

[30]Washington Business, August 26, 1991, p. 7.

[31]*Intelligent Software Strategies.* Volume X, No. 1, January 1994, p. 3.

[32]Ibid. p. 6.

This market includes not only the vendor market just summated, but also $90 million for AI application sales (off-the-shelf systems), $21 million for AI component sales (AI embedded or integrated into other systems), $244 million for AI development and maintenance costs (money paid, mostly by corporations and government, for internal or external groups that provide training, development, or maintenance of AI systems), and $45 million for Government-funded AI research conducted by commercial firms.

The lion's share of this overall market is represented by KBS systems that totaled $350 million. For KBS systems, the AI development and maintenance (or corporate) market was the single largest at $171 million, followed by the AI tools market at $141 million. Neural network and fuzzy logic systems were lumped together for this exercise, and totaled $150 million, while natural language systems amounted to $64 million. In contrast to KBS systems, neural nets and fuzzy logic, and natural language have greater relative amounts in AI application sales as off-the-shelf products than as AI tool sales. Most neural network companies have evolved from tool sales to selling off-the-shelf "task" applications, sometimes on hardwired neural network chips. Another category of AI product, logic or declarative program languages/environments, totaled $17 million.

The newsletter segments the KBS tools market into six parts. Three of the segments are based on the platform type on which the tools are generally used, and the others are separate designations for LISP tools, problem or domain-specific tools (such as diagnostic, help desk, and scheduling tools), and a rapidly growing area known as case-based reasoning tools. Case-based reasoning (reasoning based on "analogy") could be characterized as a problem/domain-specific tool, for instance, as a help desk, where it was lodged before the newsletter split it out into a separate category.

The newsletter cautions it is becoming increasingly difficult to differentiate AI revenue from revenue from the vendors' other business activities because so few "pure AI" vendors remain. Also, many vendors are seeking to distance themselves from their AI roots by shifting their marketing strategies to encompass other computing areas. These vendors are marketing their tools as "object-oriented development environments," or as more "generic computer assisted software engineering" products, or (still others) as "client-server development products." This repositioning away from the purely expert systems market niche may at first seem appealing because it allows vendors to move into larger markets and out of the currently "saturated" KBS market. However, such a move also puts vendors right into larger, even more competitive markets. Instead of having at most only a dozen competitors, vendors are now faced with large numbers of competitors from a variety of backgrounds.[33]

Table 14 presents the market statistics for each market segment for 1988-1993. The segments reporting the most growth are systems used on workstations (up 122 percent) and domain-specific tools (up 1533 percent). While the mainframe market shows growth, it actually peaked in 1990 at $43 million, but since has come down and moved sideways. In 1990, the mainframe systems market led all others, but quickly relinquished that distinction to workstations to following year.

LISP tools continue losing out to other languages, which customers have found easier to use. LISP tools are still favored somewhat by universities and research labs. The markedly downward trend in the number and value of PC and Mac tools is a sign of market maturity. Customers used these simpler tools (as "rabbits") to introduce their employees to the technology. This introductory function, however, has declined. The current small tools market is holding steady. The tools are used for such things as automating repair manuals, forms completion, and teaching aids. They perhaps are getting a little more sophisticated lately as prices rose in 1993: the number of units sold fell, while the total revenues rose.

---

[33] *Intelligent Software Strategies*. Volume IX, No. 6, June 1993, p. 1.

Table 14. Artificial Intelligence Tools Market							
Market Category	1988	1989	1990	1991	1992	1993	% change '88-'93
Knowledge-Based Systems (in millions of dollars)							
Mainframes	$11.0	$22.0	$43.0	$28.7	$29.7	$32.0	+190%
Workstations	$22.5	$23.5	$38.4	$44.0	$45.0	$49.9	+122%
LISP tools	$20.0	$22.5	$26.0	$9.6	$9.4	$9.7	(52%)
Domain-specific	$3.0	$6.0	$13.7	$21.0	$41.3	$49.0	+1533%
PC and Mac	$11.0	$6.5	$1.4	$1.8	$2.3	$2.9	(74%)
Case-Based Reasoning						$6.1	na
(in units)							
Mainframes	175	240	405	250	300	400	+129%
Workstations	3,500	3,500	6,900	7,000	7,130	8,000	+129%
LISP tools	900	970	880	460	400	835	(7%)
Domain-specific	100	200	350	790	1,870	2,200	+2100%
PC and Mac	50,000	40,000	8,420	8,700	9,080	7,380	(85%)
Case-Based Reasoning						850	na
Summary Statistics (in millions of dollars)							
Total Knowledge Based Systems	$67.5	$80.5	$122.5	$105.1	$127.7	$149.6	+113%
		Neural Net Development Tools				$25.9	
		Fuzzy Logic Dev. Tools				$6.2	
		Natural Language Dev. Tools				$25.6	
		TOTAL VENDOR MARKET				$207.3	

Source: *Intelligent Software Strategies*. Volume X, No. 2, February, 1993.

4.5.1 Vendor Survey - More than half of the larger AI vendors (over $5 million in revenues) and many smaller ones did not respond to the Department of Commerce survey questionnaire. Many of the non-respondents simply completed the block on the first page of the survey disclaiming involvement in artificial intelligence or related products. This speaks loudly to the current confusion in the marketplace and a desire on the part of these vendors to disassociate themselves from the technology. However, responses were still obtained from over 50 companies. Some of these surveys were incomplete, and certain sections could not be used.

Companies with fewer than 50 employees were only asked to provide statistical information for the calendar year 1991. Larger firms were asked to provide statistics for 1989-1992. A compilation of this information is presented on the table below. Since the statistics are incomplete, they are used only as a "performance indicator" from which inferences about the greater industry may in some cases be drawn.

With some notable individual company exceptions, the AI sector as a whole has not been profitable.[34] The group as a whole showed profits in only one year, 1990. Each of the larger firms reporting for multiple years showed at least one year of losses, and some showed losses for multiple years. One major firm showed a major drop in revenues, which if removed from the set, would result in a strong upward trend in shipments. In 1991, most larger firms showed losses, presumably because of the collapsing mainframe market, and the onset of recession. Three new firms entered the market during the years 1991 and 1992.

As a group, both large and small firms have proportionately roughly the same number of knowledge engineers and software developers, which may mean economies of scale do not pertain, at least not technically. Sales per employee is greater for smaller firms. However, larger firms can be assumed to have more consultation and training fees that were not collected in the survey.

---

[34]The profitability and financial standing of companies not completing the survey can be assumed as generally worse than those that did. This assumption is based on the idea that firms not doing well financially are less inclined and less able to consign the time and resources to completing the survey. Also, an unknown number of AI vendors went out of business during the period from whom no response was received, which would further worsen the results. The numbers reported on Table 15, therefore, probably present a better picture than in fact was the case.

Table 15. Artificial Intelligence Performance Indicators				
Performance Indicator	1989	1990	1991	1992
Sampling of Surveyed AI Vendors with 50 or more employees, 1989-1992				
Shipments of AI Tools ($000s)	$44,798	$49,369	$48,275	$52,158
Net Income (Loss) Before Tax ($000s)	($727)	$3,857	($1,431)	($2,957)
Total # of Firms Reporting Income or Loss	7	7	7	10
# of Firms Reporting Loss	3	3	5	3
Total Employment	707	741	883	970
Employment of Knowledge Engineers	116	135	160	197
Employment of Software Developers	294	290	317	325
Current Ratio (Current Assets/Current Liab.)	1.5	1.5	2.0	2.0
Total Asset Turnover (Net Sales/Assets)	1.5	1.5	1.3	1.4
Current Assets to Total Assets	64%	68%	75%	75%
Sampling of Surveyed AI Vendors with fewer than 50 employees, 1991 only				
Shipments of AI Tools ($000s)			$29,728	
Net Income (Loss) Before Tax ($000s)			($5,581)	
Total # of Firms Reporting Income or Loss			25	
# of Firms Reporting Loss			11	
Total Employment			424	
Employment of Knowledge Engineers			76	
Employment of Software Developers			169	
Current Ratio (Current Assets/Current Liab.)			.8	
Total Asset Turnover (Net Sales/Assets)			.7	
Current Assets to Total Assets			30%	

Source: U.S. Dept. of Commerce, BXA/OIRA AI Sector Survey

Larger firms get much more revenue per dollar of fixed assets, as shown by a total asset turnover rate about twice that of smaller firms. Part of the reason may be the relative complexity and expense of their respective product offerings. Larger firms tend to build larger and more expensive AI systems with an asset base not particularly different from that of smaller firms. The revenue received on more expensive systems appears to rise faster than the time devoted to their development. Thus, the more expensive systems are more labor intensive, requiring additional special customer support and follow-up services.

4.5.2 Publicly Held Vendors - To supplement the survey data, the financial statements and other materials of four publicly held AI companies were reviewed - Trinzic Corporation, IntelliCorp, Kurzweil Applied Intelligence, and Cimflex-Teknowledge.

Table 16.
Sales and Net Income of Four Publicly Held AI-Related Companies
(in thousands of dollars)

Year	Trinzic		IntelliCorp		Kurzweil		Cimflex-Teknowledge	
	Sales	Net Inc.	Sales	Net Inc.	Sales	Net Inc.	Sales	Net Inc.
1985	6,309	228	10,909	950	na	-	na	-
1986	4,980	(3,100)	18,597	4,906	na	-	14,563	141
1987	9,505	650	20,352	3,987	na	-	20,459	(2,974)
1988	20,185	122	20,433	(1,474)	1,877	(6,042)	14,383	(4,671)
1989	29,188	(1,351)	22,017	968	4,968	(3,913)	39,511	(7,833)
1990	39,904	(2,086)	19,540	(4,048)	6,769	(2,894)	25,496	(19,557)
1991	39,752	(11,730)	13,918	(11,628)	10,186	(547)	14,672	(1,151)
1992	41,344	(8,892)	9,194	(8,987)	13,880	76	7,541	(5,044)
1993	44,059	1,708	10,279	(4,772)	17,669	865	5,844	(3,114)

na - not available

Notes: Trinzic's accounting year ends March 31; on the table the company's reported "annual" figures are posted in the prior calendar year to more accurately align the data with the greater portion of business done in that year. Also, Trinzic's figures beginning in 1988 reflect the merged figures of both AICorp and AION. Prior to 1988, the figures are for AICorp only. IntelliCorp's accounting year ends June 30. Kurzweil's accounting year ended December 31 until 1991; 1992 reflects 13 months ending January 31, 1993. Trinzic's and Kurzweil's 1993 figures are annualized from 9 months of data. Cimflex-Teknowledge figures prior to the 1989 merger of Cimflex with Teknowledge reflect only Teknowledge. The company's 1993 figures are annualized based on 6 months results.

Source: Company Annual Reports and Financial Statements

Details of these companies are presented to demonstrate how some companies responded to the changing computing environment in recent years. Statements are not intended as endorsements or criticisms of the companies' products. Table 16 presents the sales and net income of these companies from 1985 (if available) to 1993.

The review showed a composite trend for these companies much worse than the survey data revealed. For example, from 1988-1993, the group reported total sales revenue of $473 million, with losses totaling $106 million, or over 22 percent of revenues. Perhaps this partly explains the disenchantment with their AI roots. As a group, sales revenue plunged from a peak of $96 million in 1989, to $72 million in 1992. Note that Trinzic and Kurzweil both showed growth during this time. An additional consideration is that Cimflex-Teknowledge had substantial revenue (and losses) from non-AI related business, which would moderate the severity of the figures somewhat if adjusted. Definite improvement occurred in 1993, as three of the four showed growth and composite sales climbed to $78 million, up more than 8 percent from the previous year.

4.6 Labor Concerns

Artificial intelligence is a multi-disciplined technology that is difficult to master. The market momentum today is to bury AI techniques in the software so that the user is unaware that it is (or it's) even there, while at the same time performing many tasks that once made computing difficult. In this context, the know-how required of an AI software developer or knowledge engineer, which has always been deep, is now also getting broader and more challenging.

The OIRA survey asked respondents to report on any difficulties encountered in the availability or hiring of qualified people during the most recent five year period. While 21 firms and 11 universities responded, 32 firms and 15 universities did not, implying that about 60 percent of the entities did not consider availability of qualified people a problem.

However, the major AI universities reported they were experiencing some problems, although none appeared to be particularly severe. In fact, one major AI institution noted it was a very prestigious place for AI scientists to work, which might disqualify it as representative of problems elsewhere in the country. The difficulties mentioned included: 1) problems in hiring people with very specific skills, 2) fewer outstanding AI faculty candidates, 3) deficits in multi-disciplinary training, notably in cognitive and perceptual science, and in mathematical/computer modeling of neural networks, 4) little appreciation among computer scientists and engineers of underlying brain anatomy and physiology, and 5) too few Master and Doctoral level people in AI. One university added that undergraduate degree level people generally have the wrong training and require too much money to retrain.

The other AI universities reported similar problems, and added others. Several cited difficulty finding qualified graduate students, particularly in the available pool of U.S. citizen graduate students. Several also noted that many students lack an adequate math background. Others stated very few people have experience in developing expert systems, robotics, and vision, and that it is difficult to find research assistants with strong, broad undergraduate training. Also, minority applicants (women, African Americans, and Hispanics) are scarce. Another reported that foreign students often have severe language problems. Moreover, many students have not learned to work on state-of-the-art equipment with current software.

Many AI firms reported they had to broaden their recruiting area to find qualified applicants, and in many cases relocate a significant percentage of new hires. Some actively looked overseas for qualified people. In general, the firms indicated difficulty finding people that combined knowledge of AI with proficiency in other technology areas, and that understood practical application issues, along with commercial software development methods and requirements.

One large AI firm that develops AI in-house for its own use, reported the initial AI person, a knowledge engineer, was "extremely" difficult to locate. Another large in-house developer firm reported when looking to hire AI people, very few apply, and those that do are very specialized. This poses a problem to the firm, because the firm needs to build AI applications in a number of different technology areas.

Many of the firms mentioned very specific problem areas. Four reported shortages of language programmers (C, C++, LISP, InterLISP, etc.). Others reported shortages of knowledge engineers, software developers, AI researchers, various programmers, and development engineers. Also, applicants with experience in various application areas, such as manufacturing and communications, were reported in short supply. One firm noted that not enough good computer science graduates were coming from the United States.

4.6.1 Labor Training - Firms and universities were also asked to report the time it takes to train a college graduate (or equivalent) to be a knowledge engineer, software developer, or research scientist. The respondents reported training times in terms of the shortest and longest intervals (in months) they experienced (or estimated) it would take to train individuals to proficiency for each occupation. Both knowledge engineers and software developers averaged from just under a year to about two years. Research scientists take longer, averaging 2-4 years. Many AI research scientists are PhDs.

Professor Bob Smith of California State University made the hypothesis that those companies with knowledge engineers on the payroll would reveal a different training interval from those companies not employing knowledge engineers. Professor Smith believes knowledge engineers, who generally better understand how people learn, know how to present both formal and informal training programs that facilitate learning on the part of new employees. In checking Professor Smith's hypothesis, the survey sample indicated a significant difference that showed companies with knowledge engineers train employees in roughly two-thirds the time taken by other companies. The former companies also show less variation in the time intervals, which would indicate more consistency.[35]

For example, as shown on Table 17, in the training of software developers, companies with knowledge engineers (23 companies) averaged 8.7 to 18.9 months, while those that did not (also 23 companies) average 12.1 to 27.4 months. The standard deviations for these averages were also significantly less for those companies with knowledge engineers. Based on the fact that companies with knowledge engineers reported significantly shorter training intervals, it can be inferred a competitive advantage is derived in this respect by simply having knowledge engineers on the payroll.

---

[35]This was put to the t-distribution test for software developers. The difference between the two sample means can be said to be true of the general population of AI vendors with 95 percent confidence. This applied to the short and long intervals, and the difference between the intervals.

Table 17. Training Schedules Important AI Occupations									
Measure	Knowledge Engineers			Software Developers			Research Scientists		
	training time in months			training time in months			training time in months		
	short	long	interval	short	long	interval	short	long	interval
Companies with Knowledge Engineers on the Payroll									
average	10.45	20.32	9.86	8.70	18.91	10.22	20.16	39.89	19.74
std	9.82	13.66	5.83	6.74	10.42	7.00	15.68	25.05	15.46
# reports	22			23			19		
Companies without Knowledge Engineers on the Payroll									
average	9.07	21.80	12.73	12.09	27.39	15.30	24.46	48.00	23.54
std	4.82	11.20	9.98	9.15	18.29	10.96	16.04	21.05	15.34
# reports	15			23			13		
Universities									
average	9.12	22.94	13.82	9.26	22.37	13.11	27.24	50.81	23.57
std	6.11	8.79	7.46	5.73	7.90	6.32	19.66	24.13	14.62
# reports	17			19			21		

std= standard deviation

Source: U.S. Dept. of Commerce, BXA/OIRA AI Sector Survey

The training patterns for research scientists are somewhat different, as might be expected. The differences between companies with and without knowledge engineers is still noticeable, if not as great as for software developers. Oddly, an even greater difference is found between companies with knowledge engineers and universities. Most university respondents to the survey do not have knowledge engineers on the payroll, but they were mixed with those who did, which further complicated comparisons. It could be that companies may have more one-on-one contact with trainees and provide on-the-job training. They may also feel the pressure of time to get the individual up to speed, and be quicker to weed out slower learners. Universities are a far different environment. The comparison is instructive, but probably not valid in this case.

Most of the vendors reporting no knowledge engineers on the payroll were small, some with fewer than 5 employees. In some cases, their orientation was toward neural networks or the like, and not toward "symbolic" representation, where knowledge engineers are most needed. However, the point is that knowledge engineers are adept at understanding how humans learn and organize information, and therefore benefit their cohorts.

4.7 Market Acceptance

A major obstacle to the commercialization of almost any new technology is market awareness and acceptance. It can take a major education effort to introduce a new technology, which from a vendor's standpoint means a lot of leg work, customer consultations, and salesmanship. From a customer's standpoint it means investment in worker training and equipment. If the technology is "revolutionary," it may even entail a re-engineering of entire business organizations, and require an often unwelcome "paradigm shift" in the thinking of business management. Computer technology is such a technology; and AI is computer technology's most advanced and difficult area.

A recent article in *PC Magazine* summed up the problem when it quoted a business survey that reported 50 percent of the CEOs in the United States and Great Britain were computer illiterate.[36] With today's exposure to computers, one must conclude computer illiteracy at this level can only occur by "choice," not a lack of aptitude, which implies an unwillingness to change. The point is that

> **"upper management" support is essential to the adoption and successful commercialization of new technology.**

4.7.1 Historic Precedence - While computer technology is now about 50 years old, and has gone through several waves of major improvements, the slowness of its entry as a "strategic

[36] *PC Magazine*, "DOS for Dummy Executives," based on survey by Robert Half International. Robert Dvorak, March 15, 1994, p. 95.

tool" into the business decision process is not without precedence. For example, Eli Whitney helped popularize the concept of mechanization and interchangeable parts when he contracted with the Federal Government to make 10,000 muskets in the late 18th Century. This ushered in the machine age, eventually displacing the predominant "craftsman" technology (and culture).

> In 1798, Whitney obtained a government contract to make 10,000 muskets. He suggested that machine tools - manned by workers who did not need the highly specialized skills of gunsmiths - could produce standardized parts to exact specifications, and that any part could be used as a component of any musket. Machines - lathes, mills, drills, and forges - had to be invented to make the process a reality. Simeon North perfected interchangeable parts (1815) using hand tools, while John Hall succeeded in mechanizing the process (1820). The chief patron for this technology remained the Federal Government, which over the next 50 years proved its value. As machine tools and other equipment improved to make better firearms, the technology gradually spread to the textile industry and agricultural equipment industries. By 1860, it was firmly established as a viable technology, known as the American System. The process became the basis for driving product prices down and raising the collective standard of living.

The required paradigm shift meant abandoning one's "home workshop" for the "factory system" and all that might entail respecting changes in life style, values, and unforeseen consequences. This is a good illustration because "interchangeable parts technology" upended the status quo. But it took many years of technical advancements, and a new generation not quite so attached to the former ways, for it to take hold. Understandably, craftsmen did not lead the effort. Eventually, business organizations were re-engineered; they went from:

> **decentralized** (i.e., thousands of small craftsman workshops, often connected or adjacent the home, making one item at a time); to **centralized** (i.e., specialized factories and division of labor making identical items many at a time).

The advent of high capacity general use computers and software to run it (such as AI) has reversed that change, accelerating a new trend toward decentralization (e.g., corporate downsizing, telecommuting, concurrent engineering, etc.). This trend has highlighted the need for an information highway that is now part of the public agenda. The computer is reengineering the way things are done, and like "mechanization" in the last century, can be expected to change the status quo in unanticipated ways. One thing almost certain will be a significant increase in overall economic productivity, with AI playing an important role.[37]

Centralized control over "decision making functions" in large organizations is part of our inheritance from the past, but has become a liability; it is dying a slow death. Centralized control has become unwieldy in a world where "decisions" must be made at the local level while new information is still "hot." That is why smaller and mid-size companies are outmaneuvering some of the giant corporations. In effect,

the computer has drastically shortened the half-life of information.

This is where AI fits in - AI refines "hot" information into knowledge (i.e., a basis for action to achieve a goal given certain information) while it is still of value, and delivers it to the person(s) who can use it. AI allows a large organization to decentralize decision making and behave with the agility of a much smaller company.[38]

As already stated, when AI vendors began commercializing AI in the early 1980s, the computing world was not yet ready. Corporate computing was dominated by mainframe computers, centralized networks, procedural languages, and top-down approaches to application development.[39] The whole corporate "information systems" culture was stuck with outdated attitudes and techniques that no one could change overnight.[40] It would have taken a major re-engineering of these systems for AI to have succeeded in a big way at that time.

---

[37]It is possible that just as a factory can today be located almost anywhere, the (knowledge) worker of tomorrow will be able to physically locate anywhere, and provide remote "know-how" to almost anywhere else on the planet (i.e., the whole world will be the workplace). In the more distant future remote links through robots or virtual reality may be common, including remotely driven factories, transport vehicles and shopping malls.

[38]See *The Global Paradox*, by John Naisbit.

[39]*Intelligent Software Strategies*. Volume IX, No. 8 (August 1993), p. 7.

[40]Ibid., p. 3.

Table 18.
AI Market Acceptance

What Would Improve Demand For AI Products		What Problems are Encountered in Selling AI Products		How Can AI be Made More Attractive in Marketplace	
Methods	**#**	**Reasons**	**#**	**Actions**	**#**
Increase Market Awareness		Market Rejection		Market Improvements	
a. Customer Education	9	a. Uneducated Customers	10	a. Educate Customers	6
b. More Advertising	7	b. Reputation Problems	10	b. Improve Image	18
A Stronger Economy	11	c. Difficult Technology	9	c. Make AI Easier-Cheaper	19
Financial Incentives		d. Management Fear	9	d. Improve Product	3
a. Investment	5	A Poor Economy	5	Financial Incentives to Buy	4
b. Other	5	Customer Lacks Funds	7		
Expanded Federal Market	5				
All Other	8				

Source: U.S. Dept. of Commerce, BXA/OIRA AI Sector Survey

AI vendors tried to get corporations to change their languages and their development methodologies in order to develop and field applications of unproven value. It did not work, at least not to expectations, although some very impressive expert systems were fielded that lived up to their promise. Today the world is gaining momentum, moving away from procedural languages like Cobol and Fortran, and beginning to embrace object-oriented computing and client-server networks. This is the work of the 1990s. Once done, companies will have the infrastructure they need to begin capturing corporate knowledge and automating human decision-making in earnest.[41]

To get a better grasp on how AI fits into this equation, the surveyed AI vendors were asked a three pronged question related to the market acceptance of AI products. They were asked:

[41] Ibid.

1) to identify factors that would improve market demand for their product; 2) what problems they encounter in selling AI; and 3) how AI could be made more attractive to their customers. A total of 48 companies responded. In each case, a leading (but not the only) problem named was "customer education." The education issue was partly ignorance, but a large component was also related to the difficulty in simply understanding the technology, and how it can help. Table 18 presents the vendors' responses to these questions.

4.7.2 What would improve demand for AI products?: Sixteen vendors commented that "market awareness" was a drag on demand for their AI products. Nine vendors referred directly to customer[42] "education" respecting exactly what AI is; how it can be used; how it differs from conventional technology; and how it can be integrated with conventional technology. Seven AI vendors reported more advertising is needed to educate potential customers, parade successes before the public, and to reduce the "fear" level among corporate management. One vendor suggested compiling a list of "reference accounts" of pioneering users of new technology for use in demonstrating the utility of AI technology to the commercial sector. The vendor added that the Government has many such accounts which would be very useful in this respect. Another vendor reported he was unable to financially support advertising, but believed it would expand demand for his product. He added that education of customers can be a major cost of doing business, and is often not recoverable. Another vendor thought advertising would help eliminate the bad publicity AI has received in recent years.

Eleven AI vendors reported that the weak economy in recent years limited sales of AI products, and predicted a resurgence of the economy would strengthen demand. One firm noted that new technologies like AI are "always the first to get cut" when the economy weakens. At a more micro-economic level, 10 AI companies reported they would benefit from various forms of financial incentives, mostly to do with raising investment capital for marketing and to upgrade product offerings. It was also mentioned that potential customers could increase demand for AI products with properly placed investment tax incentives. One firm also mentioned financial assistance would be useful to expand foreign markets and increase AI exports. Another vendor reported incentives are needed for more research and development which would help improve his software.

[42]Customer is here interchangeable with "CEO" or management.

Several vendors suggested the government market could easily be enlarged with great benefits accrued to the American people as well as to the further development of AI technology. Also, properly applied AI techniques can ease the complexity in the government procurement process, resulting in increased vendor participation, reduced costs, shortened lead times, and improved public relations. One firm suggested a specific AI requirement at least be examined, and if appropriate planned into new procurement development and implementation specifications.

In the same vein, another firm thought it would be helpful to the goal of improving public service to mandate a review of major existing government programs by knowledgeable people for AI insertion. Suggestions were also made that government agencies play a more active role in publicly advertising their use and development of AI-related technologies, reporting how the technology has improved performance in both military and civilian applications. One vendor also thought a major national initiative that requires intelligent software, such as the NASA planetary mission to Mars, would increase the demand and public visibility of the technology.

Eight additional comments dealt with a variety of issues. One vendor reported the weaker dollar helped exports, but added a policy of forcing others to drop trade barriers should continue. Another vendor suggested export controls should be simplified. A third company alluded to the difficulty in dealing with clients focused on short term considerations. Next, one AI company reported the industry needs to work on adding more to AI products (i.e., more functionality, running on more platforms), and as a group needs to address the issue of how customers can use AI technology to solve real problems better than with other methods. The firm added that this is not being done currently due to a lack of resources. Another vendor reported more AI products need to run on standard commercial processing platforms and have friendlier non-machine interfaces. Finally, another vendor pointed out that demand will increase with market maturity.

4.7.3 What problems are encountered in selling AI to your customers?: Ten vendors considered customer ignorance as the major obstacle to selling more AI products. One vendor reported most CEOs do not see AI in relation to their business problems. Another vendor reported that new markets must be created by educating the customers, which is very expensive. Other vendors reported a general lack of understanding and disbelief of the benefits of AI technology.

Ten vendors reported AI has a bad reputation, and this has hurt sales. The origins of this bad reputation lie in the overselling of knowledge based systems and creation of high expectations in the marketplace. This is frequently referred to as "media hype." At least one vendor referenced the media as most responsible for the reputation. Apparently the same hype is now spreading to neural network products, where a lot of interest has been generated with little solid information to back it up. An industry observer reported hype is now also spreading to fuzzy logic. One vendor said the AI business is still suffering from overinflated promises and hype, which makes it harder for legitimate (non-magical) solutions to penetrate the marketplace. The vendor added that the mania surrounding expert systems and the subsequent demise of many "flash but no substance" companies and projects have made customers "technophobic" to substantial applications.

Closely linked with the need for customer education, seven AI vendors reported AI technology is difficult to understand, and it takes hard work and dedication on the part of many people to build a workable system. One of the vendors said "knowledge acquisition" is a serious obstacle to the adoption of AI solutions because of the variety of specialized applications that have been developed. The vendor added that while AI is inevitably hard work, a better-educated customer population would be a huge help in overcoming this problem. Also, somehow the population has gotten the idea that AI "is easy to do," and when they find it is not, tend to discount AI solutions. The vendor also mentioned that most commercial AI vendors are still producing tools for other AI-knowledgeable people. However, very few AI systems-solution businesses exist. This puts large burdens on customers.

Nine AI vendors cited "conservative or fearful management" as a major obstacle to sales. One reported customers do not appear interested in funding research and development for a product they do not understand. Further, the customers do not appear to have clear or definable objectives. Another company reported the perception of risk on the part of customers is high due to a scarcity of successful applications. He also said potential customers resist using UNIX operating systems, and do not recognize the need for workstations (vs. PCs). Many companies appear reluctant to invest in the future. Another vendor said potential customers often do not know what AI is; they do not believe it can help; and, they do not trust very small (less than 25 people) companies. One vendor also reported customers think it costs too much. Three vendors reported customers are leery or fearful, and the management information systems groups tend to be overly cautious.

Four vendors said the poor economy in recent years reduced sales. One vendor reported in the United States, the poor economic climate, lack of incentives to try new technology products, and lack of competitive drive, seem to be the three major problems. There has not been a demand base large enough to support AI development. Another firm reported the AI vendor community is directed toward satisfying academic need above commercial need and has boxed itself into a corner by attempting to provide complex products at unrealistically low prices. A vendor also alluded to the tightness of money.

Customer-funding limitations for new technologies was mentioned by seven AI vendors. One vendor reported that difficulties in selling neural network-based processing to the Department of Defense has more to do with funding limitations than technical problems. Another reported cutbacks in government funding has created a major problem. For example, customers no longer have money to sponsor AI projects, and some are fearful of losing their jobs as cutbacks grow deeper. Another company complained about long sales cycles; the need to do cost/benefit justifications for first time applicants; and, a lack of research and development funding.

Other vendors reported a variety of concerns. One vendor reported he had multiple problems selling AI products. He said AI is perceived as having a steep learning curve, and that the technology is not integrated with existing and traditional programs. The performance of AI has been a concern because even simple applications often require a great deal of processing. Another vendor reported customizing generic expert systems is a time-consuming process for vendors and customers. A vendor also reported that customers no longer want AI for the sake of AI. They want solutions to their problems. If AI works, they want it. Otherwise, it becomes a burden of supporting more things and depending on more vendors. Another vendor reported customers perceive AI as risky due to scarcity of successful applications.

4.7.4 How can AI be made more attractive in the marketplace?: Six vendors reported "education" is necessary to making AI technology more attractive. One vendor stated flatly that "education is the key." The marketplace needs to understand what AI can and cannot do. The marketplace needs to understand the substantial successes in applications of the technology. Another vendor reported corporations need more knowledge of how AI can enhance their products. And another said customers need a better explanation of AI's reasoning process, so they can understand how the systems operate.

Seventeen AI vendors reported an "improved image" would make AI more attractive in the marketplace. Ten of these vendors suggested the technology's image would improve with the widespread advertising of success stories. For example, one vendor suggested that adoption of AI applications by highly visible leadership institutions and the general publication of those applications would help. Another vendor recommended publicizing demonstrations of where AI works to solve conventional computing problems. Other vendors reported AI needed more exposure through marketing, less hype, and more visibility when used in the public sector. Another vendor reported that he would emphasize customer productivity increases through the use of AI, which would drive costs down and competitiveness up.

Nineteen vendors reported AI products need to be easier and cheaper to use. More responses were received in this category than any other. Most reported the user base needs to be expanded; the product adapted to the PC; made more user friendly; or the technology should be integrated or embedded with other software technologies. One vendor suggested that finding ways to lower training requirements for users would help. Another said a better user interface was needed. Another vendor reported when AI is integrated with other software technologies, commercial usage will improve. The vendor added that projected advances in software and hardware technologies will also improve the attractiveness of AI.

One vendor thought the issue should focus on the customer first, and on selling "solutions," and AI should be used only where appropriate. Another vendor thought acceptance of AI as a mainstream information-management concept, or one of many software packages "in the toolbox" would improve AI's attraction.

Three vendors suggested improvements in AI products or that a refocusing of the technology to more practical use would raise market appeal. One of these vendors suggested that both better performance and a broader range of applications development need to be targeted and funded. Another suggested standardization, reliability, and robustness in AI technology are needed.

Four vendors suggested providing potential customers with tax incentives to purchase AI technology would make AI more attractive. One vendor said there is always some inertial resistance to something new. Some tax incentives for corporate America to try AI would help overcome the inertia. Another vendor mentioned several other possibilities. These included: 1) ease in procurement process; 2) special funding for proof of concept work; 3) continuous encouragement from government management; 4) special contractual vehicles for critical technology insertion; and 5) investment tax credits.

5. U.S. GOVERNMENT ROLE

The Federal Government has played a preeminent role in the development and commercialization of artificial intelligence technology. The Federal Government, notably the Department of Defense, has supplied the bulk of research funds, and provided the proving ground for many first time applications. In the future, as AI techniques become more a part of mainstream software and simpler to use, non-Defense agencies can be expected to increase their use of AI technology relative to Defense Department usage. However, Defense will remain the major user for the foreseeable future. Perhaps more importantly, the Department of Defense will continue taking new AI technology out of the lab and fielding first time applications.

AI technology is not the solution to all computer problems. However, it provides multiple improvements in efficiency, productivity, and service to the public. In conversations with various Government representatives, the evidence suggests every dollar spent on AI will save many more dollars, not only for the government, but for the private sector as well. For example, *DART* (Dynamic Analysis and Replanning Tool) solved the logistical nightmare of moving the U.S. military assets to the Saudi Desert during Desert Shield/Storm. The application was developed to schedule the transportation of all U.S. personnel and materials such as vehicles, food, and ammunition from Europe to Saudi Arabia.[43] This one application alone more than offset all the money the Advanced Research Projects Agency had funneled into AI research in the last 30 years.[44]

In another example, the Department of Labor's Occupational Safety and Health Administration (OSHA) recently completed development of an expert system, the "Cadmium Biological Monitoring Advisor," or *GOCAD* for short. The system interprets biological monitoring results and prepares documents required by the Cadmium standard. *GOCAD* uses laboratory data to classify employees and to identify corrective actions required by the rule or recommended by OSHA. The system provides expert advice to doctors, required letters to workers, grouped data for industrial hygienists, and workplace check lists for employers.

---

[43]*Expert Systems Catalog of Applications*. John Durkin, Intelligent Computer Systems, Inc. Akron, Ohio, 1993, p. 321.

[44]Dr. Steve Cross, Advanced Research Projects Agency.

GOCAD, released in late May 1994, is expected to save private firms hundreds of thousands of dollars in annual compliance and administrative costs, and conceivably millions in liability and litigation costs.[45]

5.1 What Is the Government Role?

To assist the Commerce Department in clarifying the government's role in AI, a telephone/on-line computer conference call was organized with a group of AI professionals brought together by Professor Bob Smith of California State University and financed by Steelcase Corporation in October, 1993.[46]

During this conference four major areas were discussed where the government's participation was determined to be critical. These areas are as follows:

1. **Fund Research**: Sponsor funding for basic and applied research.

2. **Purchase the Product**: Develop and deploy new and existing AI techniques that a) enable the government to accomplish its mission; b) improve government efficiency and services to the public; c) demonstrate the feasibility of the technology, and d) provide an initial market for the private sector.

3. **Manage Business Environment**: Provide legal, regulatory, and educational infrastructure to foster the development and use of AI technologies.

---

[45]GOCAD uses Instant Recall's *Tailor* expert system shell. It was distributed to the Cadmium Council and is available on LaborNews, the Department of Labor's Bulletin Board System.

[46]Conference participants were Lee Bloomquist, Principal Engineer in Research at Steelcase, Incorporated; Dr. Richard Meltzer, EDS-General Motors; Dr. Brad Cox, President of Coalition for Electronic Markets, and cofounder of Stepstone, Inc.; Paul Harmon, Editor-Intelligent Software Strategies; Dr. Bill Foskett, Office of Underground Storage Tanks, U.S. Environmental Protection Agency; Professor Emeritus Bob Smith, Director, Knowledge Engineering Lab, California State University; and John Tucker, Senior Industry Analyst, Office of Industrial Resource Administration, U.S. Department of Commerce.

4. **Leverage Risk**: Promote technology transfer, venture capital market, cooperative agreements, investments in new technologies, and dual use of militarily/commercially developed technologies.

To summarize the group's comments, the Federal Government has funded research and deployed AI products rather haphazardly, without order or general guidance, and generally without a long-term goal. The development and use of AI technologies has usually been dependent on individual initiatives, and these initiatives often faced stiff resistance from disinterested management. Also, concerns are mounting that research in the AI area may now decline as priorities and budgets shift from the AI-intensive Defense Department to less AI-intensive civilian agencies.[47]

The group was also concerned about the question of the government's role in technology development and commercialization. In a market economy, short (or long) term, low-risk projects with high success probabilities will push out high-risk projects. The government's role, then, according to the group, is basically to "broker or leverage risk." This role has two aspects. First, the government should continue funding research in basic science, and second, the government needs to broker relationships between academic and business groups to facilitate technology transfer. A graphical depiction of the public and private sector's role in technology transfer based on risk and time can be represented as follows:

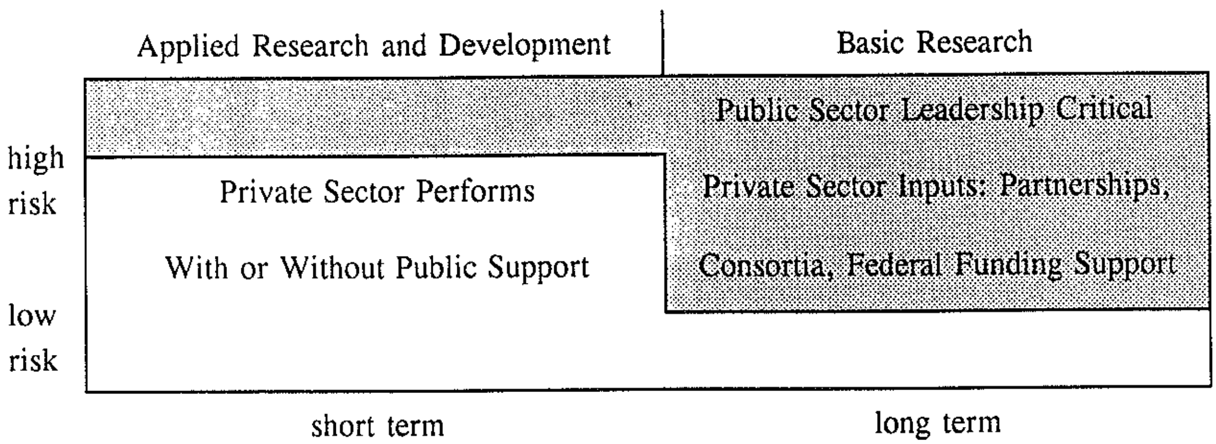

[47]White House Science Advisor John Gibbons, in a briefing February 7, 1994, restated that the goal of civilian agency vs. defense R&D funding is parity by 1998. For fiscal 1995, civilian agency R&D is projected to be $31.5 billion (44%), and Defense, $39.5 billion (56%).

5.2 Survey Comments

Companies and universities were asked to comment on the role of Federal or state governments in the development and commercialization of AI. Forty-eight firms and 28 universities responded, many with multiple suggestions. Table 20 presents the allocation of those comments into the four areas of where the Government role is critical.

Table 20.
Company and University Perception of Government's Role

Government Role	Companies	Universities
Total Number of Survey Responses	**48**	**28**
Fund Research	41	21
Purchase the Product	12	1
Manage Business Environment:		
d. Export Controls	12	
d. All Other	7	
Leverage Risk		
b. Support Cooperative Agreements	8	9
c. Change Tax Code	30	4

Source: U.S. Dept. of Commerce, BXA/OIRA AI Sector Survey

5.2.1 Universities - Universities look to the Federal Government for the bulk of their research funding. They were in general agreement that the Government needs to continue funding basic and applied AI research if the United States is to maintain its leading position. The sentiment was that the Federal government should take the lead in developing strong long-term and short-term research programs that will make the United States more competitive in the world markets. Several institutions emphasized that this is especially important to long-term (10-years or more) projects involving larger groups from industry and universities. Several also suggested funding be expanded, including more from Federal agencies other than the Department of Defense. In this regard, two schools recommended increasing funding through the National Science Foundation.

To demonstrate the importance of government seed money for research, one institution reported that neural network research suffered badly between 1972-87, because Federal support was all but terminated. This long neglect delayed maturation of this technology. Today, thanks largely to recent increased government interest, particularly from ARPA, neural networks represent one of the fastest growing AI technologies in the commercial market. The point made is that AI research areas, as demonstrated by neural networks, sometimes have a long payoff horizon and will not be funded without government support.

In addition to the Federal Government, some individual states are also financing research in selected AI areas, although in comparatively small amounts. One university reported it was currently receiving New York State support to do industrially relevant research under the New York State Science and Technology Foundation Centers for Advanced Technology. These programs have been in existence for almost 10 years. This particular university was one of nine such centers in New York State.

Universities would of course welcome increased corporate support for their AI research departments, although this may not be realistic. As things stand, more than 80 percent of total university-funded AI research currently comes from Federal and state governments, including about 75 percent of which comes from the Defense Department alone. Benefits to society could be realized more quickly with greater commercial participation in university research; however, to the commercial sector tends to support projects with "near-term" payoffs.

Collaborative arrangements between universities and private firms initiated with private funding are few in number, representing less than 10 percent of total AI research. The Defense Department funds several such collaborations, absorbing the risk by funding the projects from its own coffers. One of these, the multi-year $33 million Neural Network Program to develop autonomous vehicles, involves about 40 government, industry, and academic entities. It is believed these efforts produce fruitful results, and are excellent for transferring technology to the commercial sector.

The benefits of collaborative efforts, whether publicly or privately funded, were referred to by five universities directly, and several others indirectly. A central idea in these efforts is shared risk. Such efforts also yield faster transition and increased competitiveness. Cooperative projects involving government and private funding and university research efforts would direct research into the more productive directions (rather than shelf technology). A major AI university made the point that we, as a nation, have a need for matching real problems with funding, and added, "we could do better by promoting, through appropriate statutory, regulatory and policy changes, the use of government funded R&D consortia to transfer technology and, in particular, encourage dual use technology." Cooperative agreements would also initiate closer university/commercial ties, which might lead to additional unforeseen benefits.

Private sector resistance to such projects is due mostly to financial and competitive considerations. Providing financial leverage, for example matching funds or tax breaks, could reduce the financial disincentive. The competitive issue has in the past been alleviated somewhat by assembling firms from different industries, that do not compete head-to-head in the marketplace. A number of universities suggested supporting research through more generous tax incentives. One school indicated applied research could be expanded by providing tax credits to corporations that support university faculty. Two other institutions thought tax breaks should be targeted to long-term projects. Another recommended providing industry tax breaks targeted specifically at AI innovations.

Two other areas mentioned by universities where the government can play a role are in education and government procurement policies. The universities also expressed concern that AI may be losing some popularity among students. One suggested the government support student internships at government AI labs. Another sees a need to create incentive programs to attract graduate students into this area to build a talent pool for the future. In the procurement area, the general opinion is that procurement procedures should be streamlined and paperwork reduced.

5.2.2 Companies - The most common single comment made by the companies was for the government to fund research. This desire is all the more urgent because of the underlying need for more AI research and the companies' inability to adequately finance such research from their own earnings.

Perhaps summarizing the sentiments of the entire group of respondents, one major firm stated, "the commercial success of AI so far has been limited. Perhaps it was oversold by its early practitioners. However, there is no doubt that someday various AI technologies will mature and transform the way we relate with the computer. This is a long term view, and as such, the Government can play an important role in ensuring the United States can maintain a lead by helping to create an environment which encourages research without burdening business with undue controls."

One firm stated the Federal government should make an aggressive commitment to both basic and practical research. Federal and state governments should increase sponsorship of research and development efforts through both universities and businesses. Several companies also think more direct government involvement with financing company research would help. For example, one suggested "research grants" be offered to companies doing specific projects with good market potential. Another recommended increased investment in the Small Business Innovative Research Program, and that the gap between the award of Phase II, and the completion of the Phase I in that program be narrowed. Another company thought if the government did more research directly with small innovative companies, the high overhead costs associated with conducting research through prime contractors could be reduced.

Eight comments supported some form of government sponsorship of cooperative agreements. These agreements would link industry and academic institutions, and promote the transfer of technology to those that need it. One company indicated it is important to consider funding of ambitious multi-party projects combining industry and academic institutions, perhaps along the lines of ESPRIT-funded work in the European Community.

Two of the eight companies supported more direct government intervention. One recommended an applied AI center be established, using the Sematech model, for incubating innovative AI companies. The other firm recommended implementing an industrial policy that specifies strategic industries and improves government/industrial cooperation in those key industries.

It was also suggested the Federal government could think more strategically in the management of its research funding. For example, one company stated things would be improved if the Federal funding base for AI research (via ARPA and NSF) were more stable, and more long-term focused. Recent cuts in funding for longer term basic research, in areas like natural language processing, have left the United States vulnerable to foreign competition in the medium and longer term. Lack of stability in research funding has led to more time being spent on writing proposals than on real research. Another company added that research contracts should be based on "best value," or technical content rather than on up-front cost basis.

The Federal government could also fund research indirectly through the tax system. An R&D tax credit was instituted in the past to encourage companies to upgrade technology. Under a tax credit system, corporations are reimbursed their research expenditures by deducting them directly from taxes otherwise owed the government. A major advantage is that the R&D expenditure decisions are decentralized, and made by individual firms directly impacted by their decisions. Nine AI companies, including some of the larger players, thought this would be helpful to their operations. Another firm suggested government matching of private R&D expenditures should be considered as a more direct approach.

From the companies' perspective, the government is seen as a large potential market. However, the chief complaints about supplying the government market were delays, paperwork, and complexity in the process. A typical comment stated the process was very lengthy and complex, and often leads to a major time and financial burden, especially for small businesses. In a more positive vein, one company suggested government procurement could emphasize the use of AI. Another said AI use in the government should be expanded, while still another suggested tax laws could be changed to promote AI company interest in selling to the government. One company also said a revival of the space program would be very helpful, and another suggested the government sponsor a couple of "mega projects" with industry to increase interest and awareness.

Comments as to how the government can improve the AI business environment were the most numerous. The comments included references to export controls and other legal matters, education, and government fiscal and monetary policy.

The international market is very important to some AI companies. Government-imposed export restrictions are looked on as impediments by these companies. Twelve companies made reference to export controls, eleven of whom saw them as deleterious to their international business. The twelfth company considered the scarcity of export financing a problem where government assistance could be more most useful. In actuality, the great majority of AI software is exported without restriction to all destinations except the few nations to which nearly all trade is restricted. Most AI software is exported under general license - GTDR (General Technical Data under Restriction). Validated licenses are required for national security reasons on certain militarily significant AI software.

The problem alluded to by several companies concerns the time delays, paperwork and difficulty in understanding the export controls.[48] One company reported it had customers around the world that demand quick response. The company added that it is easy to get code (AI software products) to U.S. customers quickly, but "extremely" difficult to send code to non-U.S. customers because of export controls. Statements by several other companies supported this view indicating that the easing of export controls would make it much easier to satisfy foreign customers. One company was concerned about leveling the field by ensuring that U.S. restrictions are uniform with other industrialized countries, and that restrictions on products that have equivalent competitors in other countries be relaxed.

In the area of infrastructure, one company expressed support of the national data highway. Another company suggested the Federal Communications Commission approve a radio spectrum for computer networking. Such a spectrum would contribute to the data highway. A company also suggested "tort reform," without mention of specifics. Another firm suggested the government fund the transfer of defense-developed technology.

In the education area one firm suggested the government sponsor a major "alternative computing paradigm" conference to which significant people would be invited. The conference could include workshops, strategy sessions, and guest speakers. Another firm saw a need to stimulate demand in the commercial market with programs to educate businesses on the value of AI programs.

Tax policy, here a "leverage risk issue," received 30 comments. Several companies stated the government could help in this area by installing a more accommodating tax code. For example, the government could reinstate investment tax credits or other investment incentives. A cut in capital gains taxes was suggested by two companies. One thought special treatment of capital gains for high technology companies would generate wider diffusion of the technology. Another possibility would be to ease the tax burden on venture capital and start-up companies. The tax laws could also be used to encourage the use of expert systems and other AI products by industry, thus generating more business. Other companies suggested guaranteed loans, an expansion of Small Business Administration lend/grant programs, or direct government funding to newly formed companies or companies trying to weather economic down cycles. At the state level, one company reported

[48]Several industry spokesmen believe the export controls are an excellent candidate for an expert system that would expedite and simplify the process, and reduce paperwork.

Pennsylvania needs to repeal the tax on computer-related services and reinstate the debt operating loss carry forward provision.

5.3 Government Initiatives

During the 1980s the new buzz word became "competitiveness." In recent years, more government attention has focused on the area of technology and technology transfer with a new sense of urgency. Emphasis is now being applied in many areas that present both opportunities and challenges to private firms, including those involved in AI technology. President Clinton said,

> **"We cannot rely on a serendipitous application of defense technology to the private sector. We must aim directly at these new challenges and focus our efforts on the new opportunities before us, recognizing that government can play a key role helping private firms develop and profit from innovations."**

Major initiatives, which to varying degrees contain AI technology projects, include the Small Business Innovative Research Program started in 1982 - expected to award $700 million in grants during FY 1994; the Small Business Technology Transfer Pilot Program begun in 1992; the multi-agency Technology Reinvestment Program headed by ARPA ($554 million in FY 1994); the High Performance Computing and Communications Initiative started in 1991 ($1.27 billion in FY 1994); the Advanced Technology Program run by the Commerce Department's National Institute for Standards and Technology ($200 million in FY 1994; $451 million in FY 1995); and, NASA's Industry Technology Program, slated to receive $67 million (FY 1994). The establishment of formal lists of critical technologies by the Defense Science Board, the Technology Administration at the Department of Commerce, and the White House Office of Science and Technology Policy, is another indicator of increased government involvement, and is the basis of this study.

To maintain the competitiveness of U.S. high technology firms, the Clinton Administration is investing in a National Information Infrastructure (NII). AI and related software technologies will be essential to the operation and maintenance of NII. Designed to create a national information superhighway, this initiative seeks to establish a partnership between government and industry in which the private sector builds, operates, and improves the infrastructure and the public sector creates an efficient legal and regulatory environment and funds specific interconnection projects. Specifically, the NII has five initiatives:

1) implement the high-performance computing and communications program established in 1991;

2) create a task force on information infrastructure to implement policy changes related to the NII;

3) create a program to help industry develop advanced computing and networking technology in manufacturing, health care, life-long learning, and libraries;

4) provide funds for pilot networking projects; and,

5) promote dissemination of federal information.[49]

The Small Business Innovative Research Program (SBIR): SBIR, established under the Small Business Innovation Development Act of 1982, has been very successful at providing small businesses with opportunities to compete for federal research and development awards (including a dozen or so AI companies) and has stimulated commercialization of resulting technology. Eleven federal agencies, including Defense, NASA, and the National Science Foundation, participate in the Program, and provide (at least) 1.25 percent of their R&D budgets to small businesses. This percentage will rise gradually to 2.5 percent by 1997.

The High Performance Computing and Communication Initiative (HPCC): HPCC received $657 million in Federal funds during fiscal 1992, $964 million in 1993, and will receive about $1.27 billion in fiscal 1994. It has become the centerpiece of the U.S. Government's computer research effort. A substantial amount of funding will be used for development of generic software technology and algorithms that will help scientists address certain "Grand Challenge" problems, such as modeling changes in the global climate, finding a cure for cancer, minimizing air pollution, and improving energy conservation. Other funds will support development of scalable, parallel computer systems with a sustainable performance of one trillion floating point operations per second (teraflops) and one billion bits per second (gigabit per second) for the National Research and Education Network.[50]

---

[49]U.S. Department of Commerce, International Trade Administration, *U.S. Industrial Outlook*, 1994, Chapter 27-3.

[50]Ibid., Chapter 26-3.

Advanced Technology Program (ATP): ATP was authorized by the Omnibus Trade Act of 1988 to stimulate economic growth, create jobs, and help American companies become competitive. ATP is administered by the National Institute of Standards and Technology (NIST), part of the U.S. Department of Commerce's Technology Administration. ATP provides support on a cost sharing basis to industrial R&D projects with a significant potential for stimulating economic growth. The first appropriation for the program was for only $10 million in fiscal 1990. However, in 1994, the program was stepped up to $200 million, and will jump to $451 million in 1995. By 1997, current planning calls for the program to reach at $750 million.

Table 19.
Advanced Technology Project Commitments Involving AI Technology
(in thousands of dollars)

Company or Group	Date Authorized and Duration	Total Request	Private Funds	Major Focus
Perceptron, Inc.	Nov/1993 for 2 yrs.	$1,219	$865	Fast 3-D Image Processing
Kurzweil Applied Intelligence, Inc.	Nov/1993 for 3 yrs.	$1,777	$622	Spoken Language User Interface
Communication Intelligence Corporation	Nov/1993 for 2 yrs.	$1,480	$911	Handwriting-Based User Interface
Mathematical Technologies, Inc.	Dec/1992 for 2 yrs.	$997	$136	Mathematical Algorithms and Software
Transitions Research Corp.	Apr/1992 for 18 mos.	$699	$875	Self Navigating Robots
National Center for Manufacturing Sciences	Apr/1992 for 5 yrs.	$19,750	$26,059	Rapid Response Manufacturing
Honeywell-Sheldahl-Hercules Aerospace-3M	Apr/1992 for 3 yrs.	$2,354	$2,450	Neural Network/Sensor Control System
Communication Intelligence Corporation	Mar/1991 for 2 yrs.	$1,264	$912	Handwriting Based User Interface

Source: U.S. Dept. of Commerce, NIST-Advanced Technology Program

The focus of ATP is on transitioning technology into applications. The program funds single firms up to $2 million for (up to) three years. However, joint ventures and group projects are funded up to five years, with no specified limit placed on the amount. This provision was designed to encourage more cooperative arrangements. Through November 1993, a total of 89 projects were funded, including eight that applied to (or included) AI technology. Several AI related firms are active in the ATP program. The first was Communication Intelligence Corporation (CIC), which initiated its program in March 1991. A listing of AI projects is presented on Table 19.

The Technology Reinvestment Project (TRP): TRP has the mission to stimulate the transition (of the economy in the post Cold War era) to a "growing, integrated, national capability which provides the most advanced affordable, military systems and the most competitive commercial products." TRP programs are structured to expand high-quality employment opportunities in commercial and dual-use U.S. industries and demonstrably enhance U.S. competitiveness. This will be accomplished through the application of defense and commercial resources to develop dual-use technologies, manufacturing and technology assistance to small firms, and education and training programs that enhance U.S. manufacturing skills and target displaced defense industry workers.

TRP was authorized under the Defense Conversion, Reinvestment, and Transition Assistance Act of 1993, and other legislation. TRP was funded for $472 million in fiscal 1993, $554 million in 1994, and is slated to receive $625 million in fiscal 1995. The program is administered by the Defense Technology Conversion Council (DTCC), which includes representation from the Departments of Defense (ARPA), Energy (Defense Programs), and Commerce (NIST), the National Science Foundation, and the National Aeronautics and Space Administration. The DTCC is chaired by ARPA.

5.4 Government Agency Involvement in AI Technology

Statistics on the use of AI in the government exist only on an fragmented basis. Comprehensive data is not available, and would be difficult to gather because of the many different ways in which the technology is packaged and deployed. However, the government is unmistakably the largest single user of the technology. While data is not available, the important fact is that the government uses a great deal of the technology, and in so doing has contributed steadily to its development as a commercial product. Especially important, the

Department of Defense has on numerous occasions been the "first user" of new AI techniques, demonstrating its feasibility.

Some government agencies, or bureaus within those agencies, have formalized "AI groups" that understand the technology and actively promote the use of AI within their areas. The U.S. Army has the largest such group by far. The other armed services were slower to adopt the concept, but they now too have such groups. However, none are as large, well organized, or funded to the degree of the Army unit. Agencies such as the Internal Revenue Service, Social Security Administration, National Library of Medicine, and FBI to name a few have also adopted the specialized group approach. But most government offices are far less active, and many not even remotely aware the technology is available. In this regard, a very critical function of "institutionalized" AI groups is the training and education they provide within their jurisdictions.

The Army Artificial Intelligence Center (AAIC) is located in the Pentagon, and manages about 33 Knowledge Engineering Groups (KEGs) located all over the United States at major Army commands. Employing several hundred people, the AIC and KEGs probably have more people in AI than the rest of government combined. The mission of the AAIC is to be the Army proponent for AI applications; apply "proven" AI techniques, procedures, and methodologies at the highest levels of Army management; and, in partnership with the KEGs transfer and infuse the technology into Army activities around the world. Attached directly to Department of Army Headquarters, the AIC has been a major success not only in its stated mission, but in making the Army a more efficient and formidable force.

The AIC was established in 1984 at the direction of General Maxwell Thurmond, who recognized AI as an emerging technology and wanted it diffused into the Army. He proceeded to "institutionalize" the specialized group concept to ensure it would carry on after his tenure. His decision ultimately made all the difference and underscores the critical need for a high-level "champion" if AI is to succeed during its introduction phase into an organization. Patrick Winston, Director of the MIT AI Lab, lavished special praise on the Army unit as exemplary in its organization and effectiveness.

While the AAIC has many ongoing projects, one particularly impressive major five year and AI project under development at the Pentagon is called "***Blacksmith.***" Blacksmith is a management decision tool that will enable Generals to "simulate" with high precision how a change in policy, however large or small, will effect Army force structure and capabilities throughout the world, and thus allow decisions to be made with knowledge of the consequences.

When completed Blacksmith is designed to rapidly filter through more than a terabyte (trillion bits of data) of data, one of the largest databases in the world. Drawing from various existing sources, the database will contain the location, readiness, manning, equipment, and special capabilities of each unit in the Army. Blacksmith will be capable of "simulating" many different scenarios, and will use an assortment of statistical techniques to help present optimal policy guidance.

For example, hypothetically, if Fort X is closed in California, and the armored unit is to be reassigned to, say Fort Y in Idaho, a General might ask, "what will be the cost of doing that?" Blacksmith responds that the firing range at Fort Y is too short, and the repair facilities cannot handle the armor as currently equipped and staffed. The system will then provide a cost estimate for correction of those deficits that can be compared with other options. The General may also want to know the optimal solution, in which case Blacksmith may recommend (for example) sending the armored unit to Fort Z in Oklahoma, which currently needs an armored unit and has the facilities to handle it.

In addition to their AI groups, the services conduct AI research at their respective labs: the Office of Naval Research in Washington, D.C; the Army Research Office in Durham, North Carolina; and, the Air Force Office of Scientific Research at Wright Patterson Air Force Base. In 1991, the computer science panel of the Joint Directors of Laboratories, an interservice group, collaborated in a pilot program forming the "Centers for Excellence." AI is the "test case" for the concept which has no specified end date. Each service specializes in certain aspects of AI and offers training to personnel in the other services in that specialty. The Navy focuses on natural language understanding and machine learning, the Army on simulation, and the Air Force on planning.

The National Aeronautics and Space Administration has an AI group headquartered at Ames Research Center in California. The group has about 60 full-time employees and visiting scientists and engineers, of whom 22 hold a PhD degree in computer science or related fields, and an additional 25 hold other advanced degrees. Research is conducted at several NASA sites and contracted out to universities and others. The work is funded by NASA and other federal agencies, especially ARPA. The single largest sponsor is the Artificial Intelligence Program of the Information Science and Human Factors Division, a part of NASA's Office of Aeronautics, Exploration, and Technology. The group conducts research and develops systems for deployment. A good example of an AI system developed by the group is "PI-in-a-Box," which permits scientific experiments in the space shuttle to be executed more accurately and with greater speed.

The U.S. Treasury Department issued a report titled *Artificial Intelligence: The Application of Expert Systems Technology at the Treasury*, in February 1993. The report documented uses of expert systems in Treasury's law enforcement areas, financial management area, tax administration, and management areas. The study reported that the IRS has the largest and most mature AI program within Treasury. In its start-up phase (circa 1985) the forbears of the IRS AI Lab sent individuals for two-year graduate level study in AI at MIT, Carnegie Mellon, or the University of Pennsylvania. Other individuals received nine month training in contracting and project management and development. Today the AI Lab at the IRS employs about 40 analysts. Recently, the group contracted with a large private firm to provide a "rigorous" six-month training course in AI programming techniques.

The Internal Revenue Service (IRS) has come up with a number of excellent applications. However, IRS spent several years looking for the right AI tools for their "problems." The agency finally standardized on AI languages CLIPS and C++, but they also use LISP. They use a number of shells, such as "Level5" from Information Builders, Inc., and build others in-house. The agency also tried embedding AI in systems, with some success in case-based reasoning (reasoning by analogy), machine-learning techniques and text generation technology. The machine learning approach did not work for their application. Also, IRS officials favor in-house development of AI systems over contracting - in part because of the lengthy procurement process. Another reason is that contractors take too long to learn the complexities of tax administration.[51]

The U.S. Customs Service is working with a private contractor to develop an expert system to assist inspectors in targeting potential violators of Customs import laws and regulations. The Financial Crimes Enforcement Network (FinCen) has an AI group of 6 that is now engaged in a major effort to utilize AI techniques to proactively target possible violators of Bank Secrecy laws, using information contained in currency transaction reporting forms from financial institutions, along with financial information from several other commercial and enforcement databases to which FinCen has access. On the order of 50-60 million information bits are reviewed annually.

---

[51]Ibid.

The Financial Management Service introduced an expert system in 1989, which uses cash flow information to identify and explain the most appropriate mechanism for collecting or disbursing funds. The Office of the Comptroller of the Currency (OCC) has developed expert systems to assist bank examiners in monitoring the financial condition of individual financial institutions and to analyze their interest rate margins. Internal Revenue Service is developing an expert system which will assist tax examiners by analyzing returns to determine if an audit is required, identifying specific issues, selecting the type of audit most appropriate, and providing supporting data for the auditor.

The Environmental Protection Agency (EPA) staff introduced expert systems technology by conducting an agency-wide survey to identify the most important problems that lent themselves to expert systems solutions. Since the survey, the EPA has fielded or begun to develop dozens of applications, primarily to assist in toxic waste management. One application interprets regulatory language governing storage tanks for toxic chemicals. Another uses the technology to calculate the cost and time required to clean up hazardous waste sites. The *Expert Disclosure Analysis and Avoidance System* helps public information specialists to determine which information concerning the manufacture and distribution of toxic chemicals may be released to the public without compromising sensitive data that the EPA by law must keep from public disclosure.

The National Forest Service is producing a number of advisory systems to help foresters make management decisions. For example, the *Red Pine Advisor* is used in eastern Michigan as a forest administration aid that recommends when to plant and when to harvest. Other systems provide guidance for gypsy moth control or for the use of insecticides with other pests. In future applications, the Forest Service hopes to combine expert systems with geographic information systems to improve fire management.

In 1988 the Office of Management and Budget created an AI group to improve the management, regulatory, and budget information systems throughout the agency. Manpower reductions and the impending retirement of critical staff prompted the agency to establish an expert systems office to leverage the technological advantages offered by expert systems. This unit identifies suitable applications and disseminates the technology throughout the Office. It also provides technical assistance to users, maintains liaison with industry representatives and other federal agencies, and provides policy and technical assistance on use of the technology to the Executive Office of the President.

The Social Security Administration (SSA) has an AI group staff of 14 located in Baltimore.

One very significant program of developed by the SSA group is the "800 Expert System" or simply "800." The 800 was developed using an dated version of the "1st Class" tool, but is switching to Cincom's "ExpertRule" which is updated and easier to use. The 800 is used to answer questions from callers all over the country about the rules and procedures in obtaining SSA entitlements. The system will assist about 1,500 SSA telephone operators in responding faster and more accurately to caller's questions. In its first year of pilot operation, 800 assisted in 80 million calls; increased accuracy from 65 to 98 percent; and reduced waiting time from about 1.5 hours to a few minutes.

Sometimes too much is made of the success stories of AI technology. Very little, on the other hand, is known or documented about the failures. But there have been many of these also. While AI is gradually becoming simpler to use, especially in smaller projects, problems remain. Few, however, are actually the fault of the technology itself. In many cases, the wrong technology was applied to the wrong problem. In others, the project is insufficiently funded, or funds are withdrawn prematurely. Some are also poorly planned. In still other instances, organizational cultures are at fault.

For example, the Defense Logistics Agency was developing an "Inventory Management Assistant" expert system to help optimize inventory levels of small, inexpensive parts (e.g., nuts, bolts and screws). These items are normally managed by entry-level personnel with little experience. The AI system was designed to moderate the tendency to overstock these items, which was done to ensure their availability when needed. The system would also improve the overall performance level of the personnel assigned to manage these inventories. However, when this system was demonstrated to high-level management, it was noted the system did not follow prescribed policy procedures (as described in a 15-20 year old manual). Funds were withdrawn and the system subsequently scraped.

An official at DoD noted that uncertainty about budgeting and budget priorities has often precluded the development of AI systems, even where the payoff is potentially very large. In other words, without high-level support very little can get done. For example, the official said that not long ago DoD began funding 57 AI-related high return on investment projects at a cost of about $32 million. However, subsequent events resulted in budget cuts that killed the program. He added that with the austere mood, up-front funding for new projects is getting more difficult, despite the possibility of saving much larger sums down the road.

Unforeseen problems may also arise, despite high-level support and funding. For example, the FBI would like to put expert systems to work in criminal investigations, counterterrorism, riot management and drug operations. However, in the case of two large AI projects in the criminal and counter-intelligence fields, the AI systems did not fare as well as expected, because the field agents using them moved around too much and found it difficult to get technical support. A spokesman for the FBI's Technical Services Division said, "In more cases than not, AI systems are not as useable as we hoped. There isn't enough support, and data entry and manipulation are too complex." He added that smaller applications were more successful, citing a system to automate ballistics analysis and one to help detect serial crimes as examples.[52]

5.4.1 AI in Desert Shield/Storm - The following are examples of AI applications the Defense Department used in support of operations Desert Shield and Desert Storm.

Automated Airload Planning System (AALPS): AALPS is a military airlift load planner used by the Army and Air Force to maintain the aircraft's center of gravity, through evaluation of the shape and weight of each piece of cargo. AALPS was designed by SRI International using a graphical interface depicting the aircraft. Point-and-click on icon representations of helicopters, trucks, and other cargo are used to position cargo in the aircraft hold. The Army uses the system primarily at Ft. Bragg.

Computer Aided Embarkation Management System (CAEMS): CAEMS is an embarkation cargo load planner for ships. It was developed by the Logistics Information Systems Branch, HQ United States Marine Corps, and Stanley Associates of Alexandria, Virginia. CAEMS is a graphical interface linked to a MDSS2 database and has a built-in expert system. CAEMS provides a graphical image of ships' cargo decks and, through information in the database, permits rapid cargo stowage planning. It is deployed on the USMC 486 ruggedized PC. The system prepares a final cargo manifest with stow location for the ship's master and port of debarkation. It was deployed at Moorhead City, NC and Willington, NC.

Expert Missile Maintenance Aid (EMMA): EMMA is a diagnostic aid for the Guided Bomb Unit - 15 (GBU-15) deployed and used in Saudi Arabia in conjunction with the GJM-55 test set. The Eglin AFB Armaments Laboratory developed EMMA as part of a Productivity, Reliability, Availability, and Maintainability (PRAM) project.

---

[52] "Government Computer News." August 2, 1993, p. 8.

Pulse Radar Intelligent Diagnostic Environment (PRIDE): PRIDE is an automated diagnostic tool for the Hawk missile system's Pulse Acquisition Radar. The U.S. Army Ordnance Missile and Munitions Center, the Human Engineering Laboratory, and Carnegie Group, Inc. co-developed PRIDE. PRIDE is an expert system that combines manual procedures with the field experience of senior repairers. It provides on-screen instructional assistance to aid maintenance personnel in troubleshooting the radar. It was developed, tested and delivered to the field on ruggedized IBM compatible laptop computers in six months. It was in operation in two HAWK units deployed to Southwest Asia as part of Operation Desert Storm.

Tactical Operation Planning Support System (TOPSS): TOPSS is an operations and logistics planning tool developed by the United States Military Academy at West Point with support from the Center for Army Tactics at the Command and General Staff College. It is deployed on a Symbolics LISP machine. TOPPS provides an intuitive, graphical communications medium which promotes planning interaction between operational and logistical considerations.

The Automated Container Offering System (TACOS): TACOS is a cargo booking assistant developed the Military Traffic Management Command and Idaho National Engineering Laboratory. TACOS utilizes an expert system to automate the selection of ocean routing, port and booking of DoD overseas container shipments with the ocean carriers at lowest overall costs required to meet delivery dates.

6. DEFENSE CONVERSION, DUAL USE, AND DEFENSE SPENDING CUTS

6.1 Defense Conversion and Dual Use

Forty-one out of 43 survey respondents reported AI technology is readily transferable from defense to commercial use. The underlying AI technology, like other computer software technologies, is generic with wide applicability in its use. Many survey respondents reported that some, usually minor, modifications would be necessary for various specific AI technologies to transfer. However, a smaller group stated the change could be made cleanly. One major AI firm that performs a great deal of defense-related research reported almost all AI technology it develops for the Department of Defense is dual-use. In fact, the two firms reporting negatively qualified their answers by noting the underlying techniques would transfer, but their specific applications would require modification before commercial use would be feasible. As a result, we must conclude that

artificial intelligence is a dual-use technology.

The firms mentioned numerous dual-use areas. One firm reported that manufacturing of metal parts using AI technology is the same for both military and commercial markets. Also, defense logistics planning work is directly applicable to commercial logistics problems. For example, a major defense logistics planning project has definite applicability in many large commercial organizations for efficient inventory, distribution and transportation planning applications. Pattern classifiers (e.g., target recognition) can be used for robotic machine vision and weather pattern interpretation. The advances made in spoken language man/machine interfaces are readily transferred.

A major defense weapons system manufacturer stated that with some extra effort and money most AI technology and applications can easily be applied to the commercial sector. For example, mission planning can be applied to combating drugs, or help to diminish traffic problems of large cities. In other areas, AI health monitoring systems can support either military or commercial equipment and processes. Nuclear bomb test monitoring can potentially be used for earthquake monitoring. Database applications can be used both in defense applications and in the commercial sector.

In addition, both security and nuclear robots are dual-use. Medical imaging, non-destructive testing, and identification of planes and helicopters from their acoustics are as well. AI's fundamental tools, libraries, and algorithms are applicable in both sectors. Many applications for satellite monitoring and rocket launch control systems, originally developed for defense, are re-used for commercial projects. In case of help-desk applications, the requirements are similar for all organizations, defense and commercial. Many, many similar functions can be identified and/or mathematically redefined.

A critically important dual-use of AI is as a management decision aid. These AI systems are built to solve common (but enormous) management problems, such as building a very complex "intelligent" data base for keeping track of economic components, running a factory, streamlining a bureaucracy, or deploying military forces. AI programs, along with other software tools, have made it possible to use vast amounts of data, and turn that into useful knowledge, in many cases in real-time, to assist decision-makers.

6.2 Government's Involvement in Defense Conversion

In a two part question, the firms were asked: 1) if they were aware of any Federal, state or local government programs to assist firms in converting defense-related operations to commercial operations; and 2) what kinds of programs would be useful. Only three of 20 firms reported they knew of beneficial programs. However, a major university reported that the ARPA Dual-use Technology Conversion Initiative is a good example of a useful high leverage strategy in this area. Two of the firms mentioned that the Small Business Innovation Research (SBIR) Program was helpful, although this program was not designed specifically for defense conversion considerations. However, one of the firms reported it was using SBIR money to modify a defense application for the commercial market. The third firm indicated that useful publicly funded software is available that could facilitate conversion, but the firm lacks the capital to invest in it.

A major computer manufacturer reported that providing AI technology software to its customers, which include most major corporations, is vital to helping the United States compete in the global economy. The firm viewed defense AI systems as generally more advanced than commercial systems. Thus, providing seed money to U.S. firms to encourage the conversion of defense-related AI applications to commercial purposes would be very helpful.

The firm added that it would also be very useful for the Federal and state governments to make a greater effort to identify and deploy AI-related research developments into useful public applications. There are many instances where defense-developed technology could be applied to solve seemingly intractable government problems.

The second part to this question dealt with the kinds of programs that would be useful in assisting firms to deal with conversion problems. Here, a plethora of comments were received. A major AI firm stated that AI work is generic and software intensive. Therefore, "conversion" of operations is only a matter of funding new commercial work to replace displaced defense work. Other comments included:

- Expand rights to commercialize.

- Give grants to study market needs.

- Institute a commercial-type ARPA, especially for risky, early technology development.

- Give grants for conversion, perhaps as a step beyond the SBIR Program.

- Provide financial assistance to modify existing defense applications into reusable, off-the-shelf product technologies.

- Create a better mechanism for alerting industry and commerce to the potential uses of AI tools.

- Establish a cooperative program modeled after the ESPRIT Program in Europe.

- Include "technology transfer" tasks in defense-related Requests-for-Proposals (RFPs).

6.3 Promotion of Dual-use AI Technology by the Department of Defense

Thirteen firms and three universities commented on Defense's promotion of AI technology. All but one respondent answered that agencies of DoD do in fact promote dual-use of AI technologies either directly, or, at least from appearances, encourage it. Four respondents reported the Defense participation in the SBIR Program continuously emphasizes dual-use. One of these firms also mentioned ARPA (and NASA) as actively promoting technology transfer. On the negative side, a firm reported that government standards, i.e., the bloated and wasteful MIL spec process, inhibit commercial companies from getting involved in dual-use opportunities.

A major AI university reported that ARPA's Dual-use Technology Conversion Initiative is a good strategy, that can ultimately reduce costs to Defense by expanding commercial volumes. Another university stated that ARPA is working more closely with the National Science Foundation to both combine resources and promote dual-use. A firm reported that in the past Defense-sponsored cooperative AI projects between AI firms and universities promoted and sped-up both technology transfer and dual-use. Another firm pointed out that the Software Engineering Institute at Carnegie-Mellon University in Pittsburgh transitions defense technology into commercial use.

Various procurement groups within the DoD environment are trying to promote and proliferate the multiple uses of given AI applications. These groups are also promoting the re-usability concept, and the use of proven Commercial-Off-The-Shelf (COTS) technologies and products in defense systems as well. One firm reported that some of DoD's funding appears less related to defense, and has already helped AI companies and end users commercially. It was also stated that many research awards are oriented for the development of fundamental AI technologies, which are useful in all markets.

6.4 Impact of Defense Cuts

The Department of Defense has long been the major patron of AI research, and the leading developer and user of AI in the world. A slowdown in the Department's AI activities could severely slow down basic and applied research (much of which would not be undertaken without Federal involvement), as well as AI development. Such a slowdown would adversely effect the long-term competitiveness of the United States.

In 1993, the Department of Defense was responsible for about 60 percent of all AI research funding at both universities and companies. AI firms and universities were asked what impact the ongoing reduction of major defense programs is having (or will have) on their individual AI operations. The survey responses were overwhelmingly negative, especially those from major AI institutions. Well over half of DoD-funded AI research is concentrated at a few major universities and firms - namely universities such as Carnegie-Mellon, MIT, and Stanford, and major R&D firms, such as the Stanford Research Institute, Inc. (SRI), Bolt, Beranek and Neuman (BBN), and Loral Defense Systems. These major AI entities, as well as many others, reported severe impacts on their operations should DoD funding continue its decline. A quick tally of the results is presented on the following table.

Impact of Reduced Department of Defense Funding on Artificial Intelligence		
Tally of Survey Responses	**Companies**	**Universities**
% Reporting No Impact	23.3%	0%
% Reporting Adverse Impact Either Has or Will Occur	76.7%	100%
% Reporting Some Adverse Impact Has Already Occurred	43.3%	33.3%

Note: total number of responses: companies=30, universities=18

Source: U.S. Department of Commerce, BXA/OIRA AI Sector Survey

6.4.1 Universities - University AI labs will absorb the brunt of Defense cutbacks in AI research. A major AI institution characterized Federally funded research as a national resource in terms of transfer of government-supported technology developments to the commercial sector. Projected downturns in DoD Research, Development, Testing, and Evaluation expenditures, in particular, will impair technology transfer efforts in AI and many other fields and, as a result, impact U.S. competitiveness. Another institution noted that since the early 1960s, a significant share of research funding for computer science research has come from ARPA. Any disruption in this area, before alternative funding sources are functioning well (e.g., from a civilian counterpart to ARPA, or the like) will hurt AI programs very badly, and perhaps permanently.

An additional large chunk of AI funding comes from other branches within DoD, or indirectly, through prime contractors. Several universities reported that these sources are also declining. One university reported it has already seen several AI technology development contracts disappear from the Defense Logistics Agency, Warner Robbins Air Force Base, and the Missile Command (MICOM) at the Redstone Arsenal in Alabama. Another noted that opportunities for research grants have declined, and no doubt will continue to do so in the near future. In a similar vein, another university reported they lost at least two grants because of slowing DoD funding. One school noted as a result of less DoD-funded research, it placed more effort on obtaining support from the National Science Foundation, the Department of Energy, the National Institute of Standards and Technology, and the National Institutes of Health, as well as industry and foundations. However, despite this effort AI research programs continue at a slower pace.

The reduction in research sponsorship by DoD also adversely effects many research programs at smaller AI universities. For example, less money will be available to support graduate students, and consequently, long-term research output will be reduced. It will also lead to a significant negative impact on some institutions' ability to conduct basic and applied research in AI. Defense (prime) contractors have also reduced their funding of university AI research funded at universities. In some universities, AI research has been truncated to the point where no substantial research can be sustained. It is also certain to reduce the level of 6.3a (Defense development funding) for prototype system development. Lastly, minimal infrastructure support will not be maintained. Consequently, Defense cuts are likely to seriously reduce or slowdown the stream of new AI technology useful to U.S. industry.

Not a single university respondent reported positive impacts resulting from Defense cutbacks. However, each institution interpreted the impact differently. A sample of the comments included as follows:

1) "(The cutbacks) will reduce the pool of potential research support sources."

2) "If we cannot replace the support. we will be essentially wiped out."

3) "Minor impact to existing contracts; significant impacts on future funding opportunities."

4) "Cuts in DoD research programs could be devastating to our effort. We are highly dependent on DoD grants for all our efforts."

5) "Funding per researcher is already low in the United States. This is an alarming development."

6) "Impact may be severe, as more than 80 percent of (our) funding comes from DoD, especially ARPA."

7) "Very significant loss, could be disastrous in the future unless research money is transferred to other agencies such as NSF and earmarked for AI research."

8) "My research funding currently comes from NSF so defense will have no direct effect. Indirectly, however, cutbacks could increase competition for NSF funds."

9) "Most of our AI research has been funded by DoD. However, some of our recent grants involving vision and robotics have been funded by NSF. Major cuts in DoD funding will certainly cause us to look for other sources of funding."

6.4.2 Companies - Among AI firms, the most commonly mentioned impact of declining defense funding will be declines in sales to Defense and to prime contractors. Twenty of the 30 firms responding noted that sales have fallen, some sharply. Another eight firms reported a negative impact on their research activities, and three, on jobs. Thirteen firms reported they were already adversely impacted. On the positive side, seven firms reported the impact would be minimal or none at all. In fact, one of these noted that its AI product sales are cost-reduction directed. For this reason, sales could increase to Defense, as well as be in ever greater demand commercially. Another firm reported that its AI product increases productivity of existing personnel and reduces the cost of maintenance equipment. Further, the firm stated that new designs incorporating its product are cheaper to maintain, and potentially cheaper to produce. However, firms reporting minimal impact have little Defense business.

A very large end-user company that develops much of its own AI tools noted that more people were approaching it in search of commercial AI funding. The firm stated it must occasionally invest time to evaluate such proposals, even though there is virtually no money for such things. On the other hand, there are also occasional opportunities to recruit scientists frustrated with the increased difficulty of obtaining funds from DoD for contract research. The firm added this might be useful for an individual company, but is probably a mixed blessing for the country as a whole. In addition, the firm reported that its AI work builds on the results of academic research and often involves collaborations with defense-contract-funded researchers to transfer the technology more cheaply and quickly. Cutbacks in government funding for AI, therefore, has an indirect and adverse effect the firm.

Firms play a key role in transferring AI technology to the commercial sector. The biggest loser, however, could be their internal (mostly applied and developmental) R&D budgets, which are frequently cut in lean times. One large AI vendor noted that tighter budgets always negatively impact new technologies, including "enabling technologies" such as AI. The firm added that risk management plays a stronger role, that management focus is forced to shift, thus reducing staff and time to evaluate and instigate new programs using AI. With AI, there can always be some degree of risk, which in time can be addressed. Tighter money means tighter scrutiny and less risk taking. The firm reported it laid off approximately 10 percent of its workforce as a result of cutbacks and terminations of major defense programs. Further, it will be reducing investments in new technology development as a result of these cutbacks and, as a result, will focus its operations more on commercial domestic and international opportunities.

A large R&D firm reported that DoD (especially ARPA) has been its largest sponsor of AI programs. Cutbacks will have a major impact. Also, DoD shifts to more applied (research) areas have had a negative impact on necessary basic research. Another firm noted that Defense cutbacks have already had a very damaging impact on revenues and will dramatically reduce the ability to perform medium or long-term research and development. A third firm noted that jobs and (internal) research and development (IR&D) funds will be scaled back as defense funds are reduced.

A firm that characterized itself as a defense company reported its workload and opportunity for new AI work has dropping. As a result, the firm's IR&D and bid and proposal (B&P) fund has been substantially cut. Another large AI vendor reported decreased funding for AI has given rise to moving AI development work to the in-house operations of defense prime contractors to keep their people busy, leaving less work for vendor firms. Other aerospace and defense primes have reduced their AI staffs, and either laid off or transferred their people into other positions.

A sampling of some of the responses includes:

1) "Fewer jobs, increased lay-offs, transfer of AI staff to other technical areas. Loss of AI talent base."

2) "We anticipate either direct or indirect reductions in funded development programs and sales of software products."

3) "No impact yet. Generally we are trying to expand our business into commercial applications."

4) "Defense cuts have not directly affected us, however, we have attempted to team with various defense contractors...cuts could affect their ability to "team" with us."

5) "It may reduce the number of potential customers."

6) "...the intelligent vehicle highway system is a likely spot for converted funds."

7) "Defense applications do not account for a significant share of the Prolog (language) market. We do not anticipate much impact on our organization."

8) "Cutbacks have already had significant impact on our product sales to aerospace and Defense contractors."

9) "Over 15% of our consulting business is in the government sector. We are currently conducting marketing campaigns to replace this work with commercial contracts over the next few years, as there is no sign that there will be funding to continue some of these important programs."

10) "Our sales have decreased by one-half to two-thirds as a result of defense cutbacks."

7. UNITED STATES INTERNATIONAL STANDING IN AI

In 1993 the United States AI market was about twice the size of the rest of the world combined. However, those knowledgeable about AI consider other countries to be incorporating AI systems at a faster rate than is occurring in the United States. Economic conditions in Europe and Japan have temporarily slowed AI commercialization in those areas.

Using the technology appears to be what matters in the competitive world. In this respect, other nations, with stronger collaborative efforts between government, industry, and academia, are doing a better job than the United States.

7.1 U.S. Standing: Survey Opinions

7.1.1 Overview - Companies and universities were asked to identify areas of AI research, production, or application where the United States was losing ground or falling behind international competitors. They were also asked to name the leader(s) in those areas identified and speculate about why the United States was behind or losing ground. A total of 34 companies and 22 universities responded to this question. A summary of their views is depicted on Table 23.

Table 23. SUMMARY OF COMPANY AND UNIVERSITY VIEWS OF U.S. INTERNATIONAL STANDING IN ARTIFICIAL INTELLIGENCE				
U.S. Standing	**Companies**	**% of Total**	**Universities**	**% of Total**
U.S. leading	3	8.8%	2	9.1%
U.S. leading, but losing ground in some areas	4	11.8%	1	4.5%
U.S. leading, but losing ground in most areas	7	20.6%	3	13.1%
U.S. lead eroding, behind in some areas	17	50.0%	13	59.1%
U.S. lead eroding, behind in more and more areas	3	8.8%	3	13.6%
Total Responses	34	100.0%	22	100.0%

Source: U.S. Dept. of Commerce, BXA\OIRA AI Sector Survey

At the moment, almost everyone in the AI community perceives that the United States is ahead of the rest of the world in AI research, but that this lead is fading. The general concern among the majority of AI people at both companies and universities is that the funding needed to maintain world-class research in the artificial intelligence field may be in jeopardy due to defense spending cutbacks. This view is founded on actual and expected declines in defense research funding at some universities, particularly for long-term projects, and a drop in interest by major corporations in AI technology as a distinct discipline. This is demonstrated by very visible declines in the size and funding of many internal corporate AI groups since the mid-1980s (in part due to overall downsizing by many of our largest companies).

In the commercialization of AI products, the general sentiment was that in the United States nearly all the development risk of fielding new commercial products is absorbed by private companies preoccupied with short-term profits. This constraint imposes at most a 3-5 year payback horizon (a 1-2 year horizon is probably more realistic) on many companies, which tends to confine them to lower risk projects. This constraint may be alleviated somewhat during extended economic growth periods, "good times," and made more severe during hard times, a situation not conducive to creating a healthy transference infrastructure. The risk of development can be reduced by firms (and universities) joining in collaborative efforts, in an environment supported by the government (i.e., legal environment, tax policies, technical support, etc.). In brief, the United States lags in getting technology into commercial use for two reasons: 1) a short-term time horizon and risk aversion; and 2) poor communications between researchers and developers.

In August 1992, the National Science Board, a policy analysis division of the National Science Foundation, issued a report warning that the United States was in danger of falling behind other major countries in economic competitiveness if both industry and government spending on non-military research did not recover from an eight-year slump. It noted that the United States was still ahead in artificial intelligence and high performance computing, but trailed Japan in advanced semiconductor devices, digital imaging technology, high-density data storage, and optoelectronics. In its report, the board recommended increased Federal funding of engineering research; a permanent R&D tax credit; more cooperation between government and the private sector in science, engineering and management education programs; and NSF support for research to greatly improve manufacturing processes.[53]

Further, the United States has not effectively duplicated the caliber of various Japanese or European-style collaborative efforts between government, academic, and industrial interests. Consortia like ESPRIT in Europe and RWC in Japan have no real American counterpart. These well-funded, long-range collaborations have yielded some worthwhile results, including a more consistent and positive understanding of AI among the participants, especially in the business world, and an enlarged and growing group of skilled AI people. Moreover, many foreign graduate students receive their AI training in the United States, Ironically, several universities reported difficulties recruiting American students into the AI field.

7.1.2 Universities - The universities reported that many areas of AI are suffering from a lack of stable long-term funding. AI is a futuristic (long-term) technology and should receive support. According to these respondents, the United States has become too focused on using AI to make real things (whether the technology is ready or not) to the detriment of fundamental research. Government support for AI is too little overall. The United States will eventually fall behind if both industry and government funding for basic research continues to decrease. The U.S. Government should take the lead in developing strong long-

[53]U.S. Dept. of Commerce, U.S. Industrial Outlook, 26-3 (1993)

term and short-term research programs. One school reported the United States is generally losing ground to Japan, Germany and France in that order. Another mentioned that Japan, in particular, is gaining ground in practical areas such as applied AI and neural nets/fuzzy control.

The United States is beginning to fall behind (or has already) international competitors, particularly Japan, in applications of AI. This applies to consumer products, the ability to integrate expert systems with conventional systems, the development of very-large knowledge bases, and business-sector investment in knowledge-based technology. Also, we are behind in the design of intelligent computers which can "reason" over large volumes of data. In addition, the United States lacks database, image, and software standardization, which is having a mixed effect. It encourages versatility but discourages large cooperative projects.

The United States is behind Japanese production of fuzzy logic hardware and software because of Japan's industrial policy, investment environment, and pragmatic acceptance of unconventional technology. Japan has an "implementation based" research community. Japan took the lead applying fuzzy systems theory to develop intelligent control systems, and is advancing rapidly with industrial and commercial applications. American institutions need to continue taking the lead in theory, but at the same time develop new ways for turning theory into applications. In the case of fuzzy logic, the U.S. peer funding community frowned on it and Defense ignored it.

Germany may surpass the United States in the near future due to U.S. emphasis on LISP, in contrast to Germany's emphasis on Prolog. The United States is also beginning to fall behind Germany in some areas of knowledge representation and reasoning. Germany is rapidly and aggressively building strength in formal reasoning methods. In intelligent control systems (neural networks, fuzzy logic, expert systems) research and intelligent data visualization, the potential for advancement over our research by German and French research groups is enhanced because they enjoy stable long-term funding.

In specialty fields of AI (e.g., neural networks, fuzzy logic), both Germany and Japan are able to move their research efforts quickly to take advantage of new theoretical ideas. This is mainly because of stronger industrial leadership (usually considered very conservative) in those countries. While corporate managements are actually extremely alert to the possibilities arising from AI research ideas in these countries, this is generally not the case in the United States. AI research in the United States needs to be channeled toward practical directions. The need for feedback from the marketplace is critical.

Other countries are manufacturing products (robots, washing machines, automobiles) that incorporate fuzzy control systems. Japan is far ahead of the United States in robotics mainly because they have industry support. Europe is beating us in "neural networks" based on ESPRIT industry/academia teams. In robotics and vision, France, United Kingdom, and Germany are rapidly catching up.

Another school said that the United States is falling behind because U.S. students do not want to study science. Also, education in general is a severe problem. Our schools (public pre-college and undergraduate) are producing fewer well-prepared, dedicated students for our graduate programs. Many PhD programs are predominantly filled by foreign students who often return to their home countries after receiving graduate student support from U.S. research funds. We as a nation have to improve U.S. education in math and science, and recognize present deficiencies as a very serious problem.

Also, there needs to be an alternative to the "peer review process" for distributions of research funds. This process suppresses innovations that are off the "mainstream" (e.g., fuzzy logic, neural networks). Peer review is appropriate for large grants. Some smaller grants should be made at the grant administrator's discretion.

7.1.3 Companies - For the most part, companies responding reported that the United States is still in the lead. However, countries such as Japan, England, France and Germany are catching up. A company stated, "The United States is tremendously deficient at transitioning technology to practical applications. Japan is very good at this and often takes ideas developed in the United States to market. Fuzzy logic is a case in point. In the United States money is spent on research that sits on the shelf."

Another company said, "We are falling behind because our government does not support the transfer of R&D technology into applications. This is especially serious with decreasing defense spending." Another company said, "The United States is falling behind in the use of AI by industry. Although we still lead in many manufacturing sectors, the industry's reluctance to use new technology products will allow international competitors to catch up with and overtake the United States." One vendor also reported it is extremely difficult to obtain small business financing. This includes sources from venture capital, private investors, large corporations and banks. "Without this capital small ventures won't happen."

Various companies reported that the United States is losing ground or is now behind in certain aspects of robotics, neural networks, expert systems, fuzzy logic, and machine learning. In machine learning, most of the key people in inductive logic programming and first order concepts of learning are in Europe and Australia, not in the United States, and there appears to be a large amount of European funding supporting multiple laboratories. The United States has already fallen behind in logic programming.

One company reported that significant work is done in the United States in parallel processing, neural nets, natural language, and visual programming. However, the competition will catch up to us in the not-too-distant future, unless steps are taken to remain ahead. We are beginning to fall behind in industrial robotics, machine natural language translation, character recognition, and information dissemination devices.

Another company added that the Japanese have been very aggressive in areas like language and speech translation, funding much larger efforts than we have here. Overall there seems to be a resistance to experimenting with more "cpu" intensive AI technologies in favor of PC and workstation based technologies and applications.

Japan is also leading in all aspects of fuzzy logic, and is very strong in robotics, in part due to lack of sustained, long-term research funding and no organizational initiatives in the United States. Fuzzy logic conflicts philosophically with the U.S. approach to control by reasoning. In Japan, strong local demand for fuzzy logic has greatly aided its development. The U.S. also appears to be losing ground to Japan in the area of robotics, where U.S. funded projects are unstable. There does not appear to be any sustained effort, funding, or articulation of achievable goals. Another company reported we are losing ground on the "hardware" side to the Japanese. The company believes this is largely a result of backward-looking tax policies.

In the manufacturing area, a company said, "Japanese and European customers are doing industrial work that is far more sophisticated than anything we see here." Another company reported the corporate mentality in the United States to show good quarterly dividends is very destructive to long-term U.S. competitiveness.

The German Government has created a German Institute of AI to promote AI technology and the Japanese government has a similar effort. In the critical area of AI and software engineering, both Japan and the European Community are investing large amounts of resources in concerted, collaborative efforts on advanced software development environments, repositories, and intelligent software assistants.

A company reported that it believes the United States leads in AI at the moment, although Japan and Israel are gaining rapidly. The big problem in America is lack of long range projects - especially in the non-academic AI community. Large commercial AI projects such as ESPRIT and Prometheus are examples of such projects.

Several commercial packages for object-oriented constraint representation have emerged into the market based on work done over several years at the European Computer Industry Research Center based in Munich. European AI vendors show positive results. Vendor packages include Cosytec's *CHIP*, ILOG's *PECOS,* and Bull's *CHARME.* The United States is missing such research consortia.

In general, the lack of serious scientific review of projects in computer science is hurting the United States badly. Our country is spending money on ill-conceived computer science programs with little hope of any commercial return. The United States is, therefore, losing ground in all areas of AI quickly. European and Pacific Rim countries are now funding both more research and commercial applications; Europe through Esprit, and Japan through the Fifth Generation and now RWC.

In addition, AI is very much related to basic research in mathematics and computer architecture. Since we tend to fall behind much of the world in our early education in the field of math, we run the risk of falling behind in the long-term, because in the future we may not have the luxury of attracting high-quality scientists from abroad to our universities and companies.

A major company reported the U.S. is not behind in AI technology at the moment. The company added that a recent survey showed that the United States is a leader in software productivity for IS (Information Systems) management and systems software. However, the country is not a leader in software productivity for military applications. Because enthusiasm for AI has waned, the possibility that we will lose ground has increased sharply. For national security reasons, if not competitive considerations, it is important that we remain at the forefront of this technology.

The United States is leading the world in AI production and applications. One company reported Japan is leading in research, adding that eventually research leads to production. Thus, U.S. interests would be better served by more government sponsorship in the research area. Another company believes the U.S. is not falling behind in R&D, but we are falling behind in commercialization due to government interference in the economy.

Table 24 details the university and company comments regarding where the United States is losing ground or behind international competitors. The categories listed down the first column of the table were garnered from the company/university comments, as opposed to being chosen in advance, giving respondents freedom to focus comments on areas of greatest concern. In evaluating the table, the following method was used based on the number of responses that applied to each category.

Guide to Evaluating Table 24	
# Responses	Evaluation Parameter
1	Inconclusive, may be true for selected or narrower aspects of listed area
2	Probably true for selected or narrower aspects of listed area
3	Probably true for most aspects of listed area
4+	True of nearly all aspects of area listed
Note: Numbers entered in narrower groupings in top portion of table should be doubled for a better fit to these evaluation parameters.	

In summary, based on this evaluation method, the United States is:

1) behind Japan in nearly all aspects of fuzzy logic;

2) probably behind in some aspects, and losing ground in most aspects, of AI research;

3) probably behind Europe in establishing consortia; and

4) behind the rest of the world in commercializing AI technologies.

Table 24. DETAILED SUMMARY OF COMPANY AND UNIVERSITY VIEWS OF U.S. INTERNATIONAL STANDING BY AI AREA AND FOREIGN MARKET WHERE THE UNITED STATES IS BEHIND OR LOSING GROUND

AI AREAS * * * AI Technology	Companies				Universities			
	Japan		Europe		Japan		Europe	
	U.S. behind	losing ground	U.S. behind	losing ground	U.S. behind	losing ground	U.S. behind	losing ground
Fuzzy Logic	5		1		3	1	1	
Neural Networks						1	1	
Expert Systems			1					
Robotics	1	1			1			1
Natural Language		1						
Machine Learning			1	1				
Logic Programming		1		1				
Database Interfaces	1							
Knowledge Rep.							1	
Reasoning							1	
Hardware		1						
Responses for which only General AI Area was indicated								
Research	2	1	1			2		2
Commercializing	4		3		3			
Education	1		1				1	
Consortia			2		1		2	
Other		1	1					
Responses for which General AI Area but no Country was indicated								
Research		2				5		
Commercializing	5	1			3			
Education	1				1			
Technology	1	2			2			

Source: U.S. Dept. of Commerce, BXA\OIRA Technology Survey

The Office of Foreign Availability (OFA), a sister office of OIRA within the Bureau of Export Administration at the Department of Commerce, provided a report for this assessment on artificial intelligence in foreign countries. Parts of the report are provided here as background for European and Japanese activities in AI. The OFA report focused on expert systems, neural networks, and fuzzy logic.[54]

7.2 AI in Europe

7.2.1 Expert Systems (Europe)

European Community Research Collaboration - Most collaborative research on expert systems in Europe is sponsored by the Commission of the European Community (CEC). Under CEC's Framework program, firms and universities in different EC nations jointly undertake R&D projects on a broad spectrum of generic and pre-competitive technologies. The CEC funds up to 50 percent of the costs of projects it sponsors. In the past decade, the Framework program has sponsored several hundred projects involving several hundred firms (start-ups, middle-market, as well as large-size firms); perhaps one to two dozen involve AI. A given project may involve a dozen or more firms and universities, and generally receives less than $5 million in CEC funding.

The CEC sponsors expert systems research and development within several Framework initiatives: BRITE (research on industrial technologies), RACE (research on telecommunications) and ESPRIT (research on information technologies).[55] Research under the BRITE initiative focuses on developing expert systems applications that enable or enhance performance of manufacturing design and processes, such as applications for textile

---

[54]A copy of the AI report can be obtained from the U.S. Department of Commerce's National Technical Information Service on telephone (703) 487-4650. The report is titled, *Foreign Industry Analysis: Artificial Intelligence*, publication # PB93-183176.

[55]European Strategic Program for Research and Development in Information Technology, est. 1984

manufacturing layout or for configuring manufacture of machine parts. RACE projects focus on expert systems for configuring and maintaining telecommunications equipment, such as digital exchanges.[56]

Most CEC-sponsored expert systems research is conducted under the ESPRIT initiative. In contrast to the domain-specific, vertical applications pursued under BRITE and RACE, most ESPRIT projects focus more broadly on knowledge engineering methods, as illustrated below. The structure of the CEC is also depicted below.

Much of the ESPRIT research is oriented toward analyzing feasibility and establishing standards, although some projects have yielded results that have been commercialized. Delphi S.A., an Italian software house, markets an ESPRIT-developed expert systems shell called "Omega." Siemens Plessey has commercialized the Expert Systems Builder, an enduser-friendly development tool.[57]

While ESPRIT projects promote cross-border collaboration on an ad hoc basis, a more permanent mechanism has been realized through the European Computer Research Center (ECRC) in Munich. Bull, Siemens, and ICL (now a Fujitsu subsidiary) established the ECRC in 1984 amid concerns over potential deterioration of Europe's position in expert systems following Japan's initiation of the Fifth Generation project. With a staff of 55, the ECRC conducts research on computer operating systems, languages, and databases at a cost of roughly $8 million annually. The ECRC's participation in ESPRIT projects covers roughly 20 percent of these costs.[58]

---

[56]Al Roth, "The Esprit Initiative: AI research in Europe", AI Expert, September 1991.

[57]Ibid.

[58]Ibid.

ORGANIZATION CHART OF ESPRIT

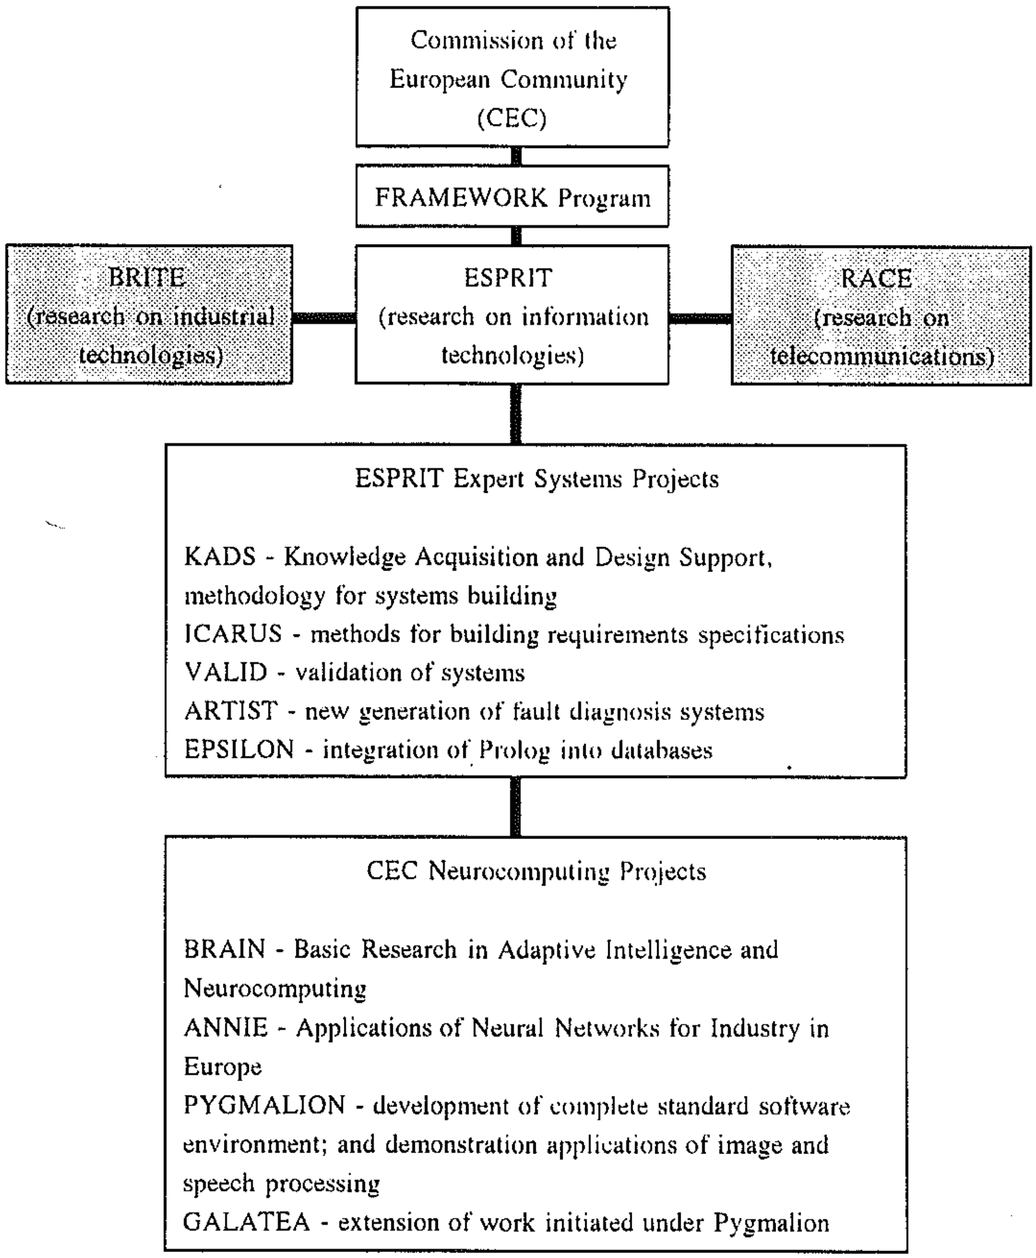

The ECRC developed a new, programming-language-based, innovative technique called constraint logic programming. The language "Chip" (for Constraint Handling in Prolog) is designed to manage complex scheduling and planning tasks. Because of its better speed, flexibility, and robustness compared with conventional expert systems methods, the constraint logic programming technique has generated strong interest among researchers internationally. Chip has recently been commercialized and introduced on the European market. Availability, however, remains limited.

The ECRC has undertaken other projects related to Prolog processing. In 1990, the center developed a Prolog co-processor chip, the "Knowledge Crunching Machine," to speed processing of Prolog-based programs on workstations. Currently, the center is researching the design of deductive databases in Prolog.

7.2.2 Neural Computing (Europe) - The CEC has sponsored four major projects to research neurocomputing. In 1987 the CEC funded the two-year, $1.6 million BRAIN project, Basic Research in Adaptive Intelligence and Neurocomputing. More than an acronym, BRAIN's name symbolized the project's intention to explore how the cognitive behavior and neurobiology of the human brain might point to new models for neural network software. BRAIN's research was basic.

Following the conclusion of BRAIN, the CEC launched two new projects in 1989: ANNIE and PYGMALION. ANNIE--Applications of Neural Networks for Industry in Europe--addressed what neurocomputing could do for industry. The three year, $6.5 million project evaluated the potential and performance of neural networks in three generic areas: pattern recognition, control,and optimization. Researchers evaluated the performance of neural networks for three applications: airline crew scheduling, manufacturing inspection, and mobile robots.[59]

While ANNIE explored the feasibility of neurocomputing, Pygmalion explored its development. Lead by Thomson-CSF, the two year $4 million Pygmalion project had two main objectives: develop a complete standard software environment (language, compilers, tools) for developing neural network algorithms and applications, and develop demonstration

[59] J.C. Collingwood, "An Overview of the first year of ESPIRT Project ANNIE," International Neural Network Conference, 1990.

applications for image and speech processing. Specific applications included remote sensing and factory inspection (image processing) and speaker-adaptive and speaker-independent word recognition for office and telecommunications equipment.[60]

Currently the CEC is funding the Galatea project, an extension of work initiated under Pygmalion. Funded at $13 million over three years, Galatea brings together 11 firms and universities to create the necessary programming tools and hardware for complete neurocomputer systems. Hardware research includes development of general-purpose neurocomputer hardware and silicon compilers for design of application-specific neurochips. Software research includes machine language compilers and graphical programming tools, as well as projects demonstrating image processing for printed circuit board assembly inspection, optical character recognition, and "videograding" citrus fruit quality.[61]

France leads Europe in neurocomputing activity. Thomson-CSF is engaged in broad-ranging research, and a half-dozen small French firms have entered the neural network software market.

At the center of both the Galatea and Pygmalion projects, Thomson has undertaken systematic research into the potential application of neural network architectures to the firm's diverse business lines--equipment such as radar, sonar, video and infrared image processors, and systems for air-traffic control, battlefield management and telecommunications. In signal and image analysis applications, neural networks offer an approach to recognizing visual and acoustic signatures, and to facilitating low-level processing of images--compression, noise reduction, and multiplexing. Although some neural network capabilities appear promising, Thomson notes the technology is still in its early stages and more evaluation will be necessary.[62]

---

[60]Bernard Angeniol, "PYGMALION - Neurocomputing," International Neural Network Conference, 1990.

[61]Roger Woolnough, "Neural-net use studied," Electronic Engineering Times, February 24, 1992, p.31.

[62]F. Vallet, "Neural Networks Activities at Thomson-CSF," International Neural Network Congress, 1990.

Of France's neurocomputing start-ups, Mimetrics has gained the highest profile. Founded by the former head of Thomson's AI group in 1991, Mimetrics has been quick to introduce products on the market. It launched a commercial version of Mimenice, a C-language neural network simulator first developed in the Pygmalion project.[63] Mimetrics also bought out Inovatic, a French firm specializing in optical character recognition (OCR), which markets ostensibly the first neural-network-based OCR software, Easy Rider. Mimetrics plans to focus on OCR, facial recognition software to control individuals' access to facilities, and programs to forecast foreign exchange.

Projecting $4 million revenues for 1992, Mimetrics states it has already achieved profitability. Outside investors include Thomson (20%), France Telecom (10%), Finovelec (15%) and Baring Venture Partners (20%). Thomson and the French Post office are major customers. Mimetrics also maintains ties with Thomson by serving as technical leader for the Galatea project (Thomson handles Galatea's administrative side).[64]

Under Galatea, two firms, Philips and Siemens, are developing neurochip designs intended eventually for commercial production. In France, a team of 20 researchers at Philips has developed a VLSI neurochip with on-chip memory named "L-Neuro" for processing small size arrays of inputs. Philips foresees that the L-Neuro will be used on add-in boards for PCs and workstations; use in consumer electronics may follow. Potential applications include image processing, games, handwriting recognition, and speech recognition.[65]

Siemens' Research Institute in Munich has designed a neurochip, the MA16, able to process larger arrays of inputs. A prototype MA16 was completed in 1992 after three years of work. The Siemens chip is appropriate for large data sets such as television camera images.[66]

---

[63]Claire Remy, "Neural Networks: Mimenice, Development of a Simulator," Zero Un Informatique, September 27, 1991, p. 16.

[64]"First Launch Products that make People Talk About You," Electronic International Hebro, June 11, 1992, p. 33.

[65]Jean-Pierre Cahier, "Awaiting the Neural Computer," Le Monde Informatique, November 4, 1991, p. 44.

[66]"Siemens starts on neural work," Electronic Times, December 5, 1991, p. 7; Hans Schmidt, "Artificial Prophets," Bild der Wissenschaft, April 1992, p. 40.

7.2.3 Fuzzy Logic (Europe) - Despite the appearance of several small national-level collaborations researching and exchanging information about fuzzy logic, no pan-European research effort has yet emerged.

In France, a number of leading industrial firms exchange information about fuzzy logic through the Research-Industry Exchange and Coordination association (ECRIN).[67] In Germany, the North Rhine Westphalia state government is spending about $3 million a year on a "Fuzzy Initiative" to educate small- and medium-size manufacturers about fuzzy logic software techniques that can be applied to production.[68] Italy's government is providing grants to the Microelectronics Research Consortium in southern Italy, CORIMME, to develop fuzzy logic applications for use in cars.[69] (CORIMME is a joint venture between researchers from the University of Catania and SGS-Thomson Microelectronics.)

The above programs highlight a growing movement in Europe to educate firms about fuzzy logic, but none involves the cross-border collaboration that exists with expert systems and neurocomputing. Interestingly, in the absence of a big-league pan-European research project on fuzzy logic, Thomson and Volkswagen have elected to join Japan's MITI-sponsored Laboratory for International Fuzzy Engineering (LIFE).[70]

The current absence of a CEC-sponsored project in fuzzy logic may be only temporary: as interest increases, firms and national governments may voice more interest in such a project. However, the condition of Europe's electronics industry has been working against cooperation.

---

[67]Henri Pradenc, "Fuzzy Logic Serving Industry," Zero Un Informatique, November 15, 1991.

[68]"North Rhine-Westphalia in the Forefront of Fuzzy," VDI Nachricten, April 17, 1992.

[69]"Artificial Intelligence Against the Risks of the Road," Il Sole-24, March 28, 1992, p. 17.

[70]Nihon Keizai Shimbun, August 13, 1992, p. 5.

Worldwide research on fuzzy logic is largely focused on creation of chips, an area where Europe has floundered. Chip research collaboration has generally been sponsored under the Joint European Submicron Silicon Initiative program, which has been strained by financial problems as troubled European chipmakers drop product-lines and withdraw from the program.

The CEC's ESPRIT program, another program experiencing money troubles, might provide an umbrella for certain fuzzy logic research, albeit one more likely directed toward generic issues.

7.3 AI in Japan

7.3.1 Expert Systems (Japan)

National Research Collaboration - In marked contrast to ESPRIT's funding of numerous mini-projects, MITI folded nearly all research collaboration on expert systems and logic programming into the Fifth Generation program. Creators of the Fifth Generation program envisioned development of a computer able to understand natural language, access vast database libraries of knowledge, and solve large scale problems. To coordinate the research, MITI established ICOT, the Institute of New Generation Computer Technology. Between 1982 and 1992, ICOT spent roughly $340 million and hosted 100 researchers from 18 organizations, who provided the bulk of ICOT's budget.

Observers both inside and outside Japan expected that MITI's initiation of such a colossal research undertaking would yield unassailable commercial dividends for ICOT's members. Ironically, the sheer ambition of the ICOT's goals forced its research to veer in largely academic directions, isolated from immediate practical application. Research focused on new parallel computer architectures, new programming languages, and new operating systems. In an effort to realize the program's expansive goals, ICOT researchers decided to try new conceptual models in each of the areas.

ICOT's pursuit of new conceptual models, especially massively parallel symbolic computation, created unforeseen problems to understand and overcome. Over time, ICOT scaled back its research agenda considerably.

> **Attention moved away from visionary goals toward more basic implementation ones. Over its decade-long effort, ICOT demonstrated some novel ideas but no breakthroughs. At the same time, computer and AI developments outside of ICOT pointed to more promising avenues of research. More importantly, ICOT's conceptual models failed to show commercial viability. Few of the specific insights gained can be easily incorporated into commercial products.**[71]

MITI is now exploring ways to salvage more from ICOT's efforts. Recently it has encouraged foreign researchers to critique ICOT's work and cooperate on further research. MITI is proposing new funding to extend ICOT's work beyond its official conclusion in 1992. Despite the general consensus worldwide that ICOT has contributed little of value to the state of AI, MITI has concluded that "it is important to extend the results of the past ten years under a new national project."[72] A five or six year extension of ICOT will study expert systems applications that can run on the specialized parallel processing hardware ICOT developed earlier.[73]

Together with the Post and Telecommunications Ministry, MITI is offering $16 million "Key Technology Center" funding to AI Language Research Institute (ALR) to adapt ICOT's work on logic programming for commercial use.[74] Researchers from Hitachi, Fujitsu, NEC, and seven other firms will work through ALR to port the ICOT-developed Extended Self-Contained Prolog (ESP) language to Unix workstations.[75]

---

[71]JTEC Workshop of Knowledge-based Systems in Japan, June 17, 1992.

[72]"MITI proposes expansion of 5th Generation Project" Nikkan Kogyo Shimbun, July 1, 1991, p. 1.

[73]"Japan: 5th Generation Computer Project Extended" Newsbytes News Network, August 19, 1991.

[74]For a general discussion of the Key Technology Center Program, see The Government Role in Civilian Technology, National Academy Press, 1992.

[75]Japanese Technical Literature Bulletin, December 1991, p. 8.

7.3.2 Neural Computing (Japan) - In the late 1980s, two government research labs started projects on neurocomputing. The Advanced Telecommunications Research (ATR) consortium, funded in part by proceeds from the privatization of the Nippon Telephone and Telegraph (NTT) Corporation, began exploring application of neurocomputing to telephony. Loosely modeled after Bell Labs, ATR has hosted researchers from NTT and NHK, among others, to investigate speech recognition issues. The Electro-Technical Laboratory (ETL), part of MITI's Agency for Industrial Science and Technology, also started work on speech recognition, as well as data compression.[76]

A major collaborative research effort devoted to neurocomputing is only now beginning. Following close-out of the Fifth Generation project, MITI has anointed neurocomputing as the leading candidate for design of intelligent computers.

MITI has crafted a new ten-year project intended to enable design of a "flexible" computer able to process in parallel a flood of unorganized, massive, and distributed information, much like how the human brain takes in visual images.

MITI has officially christened its plan, which has been known variously as the four-dimensional computer, sixth-generation computer, or new information processing technology, as the "Real World Computing" (RWC) program. The RWC program will consist of a constellation of projects addressing three research areas: optical computing, massively parallel processing, and neural systems.

Through experimentation with prototype systems for each of these three areas, and development of new theory, MITI hopes researchers will be able to combine elements to achieve prototype computer systems that imitate intuitive aspects of intelligence, rather than simply abstract logical ones.

---

[76]R.B. Davidson, "Non-U.S. Artificial Neural Network Research." SAIC, October 1991, p. III-10,11.

MITI projects that the RWC program could cost as much as $800 million over its ten-year life. Actual spending, which will involve a mix of both MITI- and private-participant funding, may be only half of that.[77] The share of research spending related to neurocomputing is not fixed, but will not necessarily be the most prominent aspect of the RWC program.

The RWC program will focus on four topics relating to neurocomputing: neural models, hardware, software, and integration with massively parallel systems. Research on new neural models will provide a basis for new designs for neurochips. The RWC program anticipates exploring digital, analog, and hybrid approaches to neurochip designs. Preliminary goals target chips with 40-60,000 neurons after the first 5 years, improving to 1 million neurons on 100 wafers at the conclusion of the ten-year program. The program further anticipates developing new languages and operating systems to work with the new chip designs.

The RWC program has sparked interest from foreign observers, but not the alarm that greeted announcement of the Fifth Generation project. The program and its results will be available to foreign researchers. The RWC research will more resemble basic research than industrial policy: planners expect that much of the knowledge developed as a result of the program may not have practical application for 10, 20, or even 30 years.

In fact, goals of the research are so general that it is difficult at this point to speculate what will be its possible impact on the state of neurocomputing.

---

[77]Spending on some prominent MITI projects, including the Fifth Generation program, has significantly fallen short of initial projections.

7.3.3 Development of Fuzzy Logic in Japan

Japan has been almost single-handedly responsible for popularizing fuzzy logic.

Japanese firms started fuzzy research programs, created fuzzy products, and hyped the fuzzy mystique to consumers. In this way, Japan created a fuzzy boom.

It was Japan that rescued an obscure mathematical theory developed by an obscure Iranian-born, American university professor, and transformed the theory into a high-profile, mass-market technology sought by gadget-loving consumers.[78] In the early 1980s, Japan showcased the application of fuzzy logic to the operational control of trains and subways and captured the attention of Japan's leading electronics companies. By the late 1980s, consumer products abounded that were pitched as having the fuzzy advantage--products like washers, TVs and camcorders.

Japanese firms have developed many pioneering applications. Product designers have shown fertile imaginations when conceiving new applications for fuzzy logic. In addition to the boundless models of appliances, fuzzy logic has been used in a variety of industrial equipment and manufacturing processes such as blast furnaces. Fuzzy algorithms can make adjustments in some control problems more accurately than traditional approaches such as proportional, integral, and differential (PID) controllers. However, Japan's attraction to fuzzy logic seems less related to its quantitative performance advantages over PIDs than its relative simplicity. It is because of this simplicity that fuzzy logic has found a home in such low-value products as hot plates.

However, simple implementations of fuzzy logic offer only limited capabilities. Japanese products feature the fuzzy label prominently, but sometimes (such as in hot plates) the marketing behind the product is more sophisticated than the technology. Simple implementations in products can respond only to a limited range of predictable situations. Some implementations do no more than look up a table of stored values.[79] Most

---

[78] An obscure academic no longer, the conceptualizer of fuzzy logic, Dr. Lofti Zadeh, now consults Japan's top fuzzy firm, Omron, in addition to teaching at Berkeley.

[79] "Japan Starts Fuzzy Logic II," Machine Design, September 10, 1992, p. 16.

implementations are software simulations running on conventional chips, rather than dedicated fuzzy chip designs. Most contain only a few rules to process, limiting functionality (and flexibility). Presently, application of fuzzy logic is problematic in "unstable" and "ill structured" situations, in multicriteria decisions, and for interpolation.

However, for each of these limitations, changes are forthcoming. The latest generation of Japanese fuzzy products are getting more sophisticated. And, researchers are tackling barriers of theory and pushing the boundaries of the technology. Curiously, the pervasive presence of fuzzy logic in Japan belies the technology's embryonic development.

National Research Projects - Japan has several consortiums probing fuzzy logic. Their activities, while not necessarily coordinated, are often inter-related.

The Fuzzy Logic Systems Institute (FLSI), founded in 1990, is the creation of Kyushu professor Takeshi Yamakawa, a prominent fuzzy researcher and designer. Receiving administrative backing from the Education and Trade and Industry ministries, Yamakawa convinced about a dozen firms to sponsor a new $5 million laboratory dedicated to fuzzy research. Yamakawa traded patent rights to his research in return for corporate funding of his lab. Initially, rights to FLSI research were assigned to individual firms, but FLSI changed this practice after Omron cornered 65 choice patents. Now FLSI holds title to patents, which are licensed to all member firms, who pay royalties on a scale inversely proportional to the amount of funding they give.

FLSI conducts research on fuzzy software and fuzzy chips. The software is intended to create more friendly interfaces between people and computers. FLSI is pursuing several directions in chip designs. Yamakawa's team has developed a "fuzzy neuron" chip, fusing fuzzy and neural technologies, which may have application for recognizing handwritten characters. FLSI researchers have also designed an analog fuzzy chip for high speed operation. FLSI has constructed a foundry to produce its chip designs.[80]

A far larger research collaboration is the Laboratory for International Fuzzy Engineering (LIFE). Approximately 50 companies, including several foreign firms, have joined LIFE, paying from $200,000 to $500,000 for membership. LIFE is researching fuzzy control,

---

[80] R. Colin Johnson, "Japan Sets Fuzzy Group," Electronic Engineering Times, September 3, 1990, p. 18; "Fuzzy Logic Flowers in Japan," IEEE Spectrum, July 1992, pp. 32-35.

fuzzy information processing, and fuzzy computing, covering basic research, development tools and demonstration applications in each of these areas. Applications under research include decision support for foreign exchange trading and home robots able to recognize images and natural language. Most recently, LIFE has embarked on a three-year fuzzy logic programming project aimed at enabling PCs and workstations to understand natural language commands.[81]

LIFE is part of Japan's national R&D program, and receives funds from the Japanese government, which has agreed in principle to supply up to half the LIFE's funding. So far, however, MITI has hesitated to allocate large funds for LIFE. Unlike FLSI, LIFE is not expected to create patentable technologies; commercialization will occur off-site by individual member firms.[82]

Japan's Science and Technology Agency (STA) also sponsors fuzzy research. It has picked up some early research developed by LIFE on fuzzy control of helicopters. The helicopter research addresses thorny issues involving fluid, unpredictable situations. Japan's Transport Ministry has also sponsored fuzzy research related to helicopter control. Research has been contracted to the Tokyo Institute of Technology and two private firms.[83]

STA, through its Japan Development and Research Corporation unit, also awarded Omron a $5 million contract to develop technology for manufacture of analog fuzzy chips. Originally, the technology was conceived by Professor Yamakawa, now head of FLSI.[84]

Despite the Japanese government's promotion of fuzzy research, its role has been more reactive than proactive. Private concerns, rather than MITI, were responsible for proposing the FLSI and LIFE consortia. Notwithstanding MITI's legendary foresight in identifying

---

[81]Masayuki Miyazawa, "Fuzzy Computing Project to Start Soon in Japan," Newsbytes, August 5, 1992.

[82]"Fuzzy Logic in Japan: Clearly Important and Becoming More So," NTIS Foreign Technology Newsletter, May 8, 1990.

[83]"Fuzzy Flight Control Systems for Helicopter Under Development in Japan," NTIS Alert Foreign Technology, January 28, 1992, pp. 2-3.

[84]"Japan's Omron Committed to Fuzzy Control Applications," NTIS Foreign Technology, November 19, 1991.

emerging technologies for companies to target, the ministry created a fuzzy study group only in 1991, years after Japan's leading firms had committed to the fuzzy market.[85]

7.4 The JTEC Report

In addition to the OFA material, the Japanese Technology Evaluation Center (JTEC), located at Loyola College in Baltimore, issued a panel report in May, 1993, titled, "*Knowledge-Based Systems In Japan.*" The report was funded by ARPA. The panel included distinguished experts in the AI area, including Ed Feigenbaum, who chaired the project. The July, 1993 *Intelligent Software Strategies* newsletter provided an overview of AI in Europe.

The following are excerpts from the JTEC executive summary.

o Japanese computer manufacturers (JCMs) play a dominant role in the technology and business of expert systems. The JCMs have mastered and absorbed expert system technology as a core competence. They tend to use systems engineers rather than knowledge engineers to build systems. Consequently, integration with conventional information technology poses no special problem for them, and is handled routinely and smoothly, without friction. These large computer companies also build many application systems for their customers; smaller firms specializing in AI software play only a minor role in applications building, compared to the United States.

o The majority of the Japanese expert systems tools are developed, sold, and applied by the JCMs. They have the resources to conduct research, develop new products, and persist in the business. In the United States, most of the expert systems tools are developed and marketed by a handful of small companies. The Japanese can continue to invest in the research and development of new tools (which they are doing) and are in a better position to survive lean times. In contrast, American vendors must work with short-term objectives and lean cash reserves.

---

[85]"Ministry Creates 'Fuzzy' Research Group," Newsbytes, March 18, 1991.

o Japan has more experience than the United States in applications of knowledge-based systems technology to heavy industry, particularly the steel and construction industries. In certain application tasks, such as closed-loop control, expert systems have been assimilated into the suite of techniques available to the systems engineers, and do not require the special attention sometimes afforded new technologies.

o The Japanese are ahead of the United States in the integration of problem solving techniques, due to a combination of factors. These include substantial Japanese investments in experimenting with a wide range of technologies and in-house development of expert systems tools by Japanese computer manufacturers and other large organizations. These factors provide the understanding necessary for full integration of software with other data processing components. Another factor is the avoidance of artificial partitions between various methodologies.

o Products based on the use of fuzzy control logic have had a big impact on consumer products, including camcorders, automobile transmissions and cruise controls, television, air conditioners, washer/dryers, and many others.

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

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

Table 21.
U.S. Standing in Knowledge Based Systems Relative to Japan

Measure	Current State	Trend
Quality of the best	even	progressing equally
Quantity relative to GDP	even	US losing ground
Support Structure	US behind	progressing equally
Tools	US behind	US losing ground
Consumer Products	US behind	US losing ground
Integration	US behind	US holding or losing ground

Table 22.
U.S. Standing in KBS Research Relative to Japan

Area of Research	Quantity		Quality	
	Current State	Trend	Current State	Trend
Advanced KBS Research-Industry: Basic Research	US ahead	progressing equally	US ahead	progressing equally
Applied R&D	even	US losing ground	US behind	progressing equally
Advanced KBS Research in Universities	US ahead	US losing ground	US ahead	US losing ground
National Initiatives: Parallel Symbolic Computation	US behind	US gaining ground	US behind	progressing equally
Very Large Knowledge Bases	US behind	US gaining ground	even	progressing equally
Fuzzy Logic Systems	US behind	US losing ground	US behind	progressing equally

Source: *Knowledge-Based Systems In Japan*, Japanese Technology Evaluation Center, May 1993, p. xii

8. CONCLUSIONS AND RECOMMENDATIONS

8.1 Conclusions

1. The Federal government is AI's most important patron and largest customer. However, a low profile and lack of leadership in this area by the government has slowed the commercialization of the technology.

2. The downsizing of Defense has resulted in a smaller share of the research dollar going to basic research. It has also resulted in cuts in total research that will not, under current policies, be made up by increases from other Federal agencies. The U.S. Navy and Air Force have both significantly cut AI research, and ARPA appears to be cutting back about 15 percent from recent years.

3. Because of capital constraints, the private sector cannot support most basic research with longer-term payoffs needed to push AI up the evolutionary ladder. Government funding of basic research is imperative to the long-run success of AI.

4. AI is a strategic asset. AI systems, large and small, often result in a ten-fold or more productivity increase. These increases are possible because knowledge (i.e., the ability to take a specific action to achieve a goal on given information) is the most underutilized asset in any organization. This often is not yet understood by managers. The best AI systems save companies (and the government) millions of dollars a year.

5. Since about 1980, corporate computing has been shifting from centrally-controlled mainframe-based to decentralized multi-platform based. Knowledge automation of widely scattered organization information is growing in importance and stimulating the market for AI technology.

6. The United States has by far the largest AI research capability in the world. Although originally highly concentrated at a few institutions, it is now spread widely across the country. Many institutions have first-class talent but lack the funds or facilities to fully support AI research.

7. Software in general, and AI in particular, is used to tap the tremendous potential of the computer. Yet software development productivity, including AI, has stagnated for over 20 years, while the demand for and complexity of software has increased at a rapid rate, especially in recent years. Since so much of our economy is now dependent on software, it has become a major constraint to economic growth. AI has the potential to automate the software development process (in much the same way as machine tools automated physical production), and also supply reusable software component parts.

8.2 Recommendations

- **Support Cooperative Agreements**: Federal leadership is needed in defining long-term goals and strategies, and establishing greater coordination of research and development efforts among industry, academia, and Federal agencies.

- **Support Increased Software Development Productivity**: AI has a role in software productivity.. U.S. Government policymakers need a clear and objective understanding of the importance software now plays in literally every economic sector, and in so many aspects of our lives. AI techniques have the potential to increase software development productivity (and quality) by automating the complex cognitive processes software programming involves. A task force of experts in this area is needed to both discuss the situation and form a plan of action. The government, perhaps with ARPA-NIST jointly playing a lead role, can sponsor this effort and provide funding to support it.

- **Support Basic AI Research**: Federal funding is needed for basic AI research. The amount and specific areas of that funding should include industry and academic perspectives under Federal leadership. Stable, long-term funded projects are needed in dollar amounts large enough to support full-time, dedicated staff.

- **Establish Regional AI Labs**: AI talent and facilities are spread thinly at many universities and corporations. Establishment of regional AI labs with permanent support staffs and computer equipment would permit better and cheaper use of this talent. These could be self-funded.

- **Establish a Federal AI Group**: Form an interagency AI group that promotes and oversees Federal AI activities. The group could help coordinate government, commercial, and military interests; stimulate the use of AI by the government; educate public servants in its use; promote greater coordination in AI research among government agencies; promote policies and regulations to assist the development and commercialization of AI; and advise the Congress and the President on the status of AI.

- **Strengthen Statistical Tracking of AI**: The U.S. Government needs to strengthen its statistical tracking of AI technology research and development expenditures, and its purchases of AI systems.

CRITICAL TECHNOLOGY ASSESSMENT OF THE U.S. ARTIFICIAL INTELLIGENCE SECTOR

APPENDIX A

GLOSSARY OF AI TERMS

U.S. DEPARTMENT OF COMMERCE
BUREAU OF EXPORT ADMINISTRATION
OFFICE OF INDUSTRIAL RESOURCE ADMINISTRATION
STRATEGIC ANALYSIS DIVISION

APPENDIX A

GLOSSARY OF AI TERMS

ARTIFICIAL INTELLIGENCE - 1. An interdisciplinary approach to understanding human intelligence that has as its common thread the computer as an experimental vehicle. This definition emphasizes the fact that many disciplines contribute to the field of artificial intelligence. These professions include computer science, psychology, mathematics, physics, philosophy, engineering, and business. 2. The aspect of computer science that is concerned with building computer systems that emulate what is commonly associated with human intelligence. 3. An area of computer science that focuses on symbolic, nonalgorithmic methods of problem solving. 4. The building of programs that are characterized by symbolic representation, symbolic inference, and heuristic search. 5. Using symbolic pattern-matching methods to describe objects, events, or processes, and to make inferences. 6. The study of mental faculties through the use of computational models. This definition ensures the inclusion of vision and natural language processing. Major areas of artificial intelligence include robotics, general problem solving, machine learning, pattern matching, pattern recognition, logic programming, theorem proving, expert systems, game playing, decision making, planning, automatic programming, intelligent computer-aided instruction, natural language processing, vision, speech recognition, search, knowledge representation, knowledge acquisition, expert data base systems, neural networks, understanding systems, and uncertainty.

AUTOMATIC PROGRAMMING - The area of artificial intelligence involved in software construction that produces programs from a user's specifications. Computer programs that can derive programs from being told what is wanted rather than how to accomplish the task.

DECISION-MODELING SYSTEM - An artificial intelligence program involved in selecting the best solution from a number of known alternatives. These programs generally use decision trees. A decision modeling system differs from an expert system in that with decision modeling software it is a matter of choosing between one of a small number of known alternatives, as in choosing one of five different stock portfolios.

EXPERT DATA BASE SYSTEMS - A combination of a data base system and an expert system, which is most applicable when used with complex and ill-structured data. They have a flexible knowledge representation scheme. See also Knowledge-base management systems.

EXPERT SYSTEMS - A computer program that uses symbolic knowledge and inference to reach conclusions. It derives most of its power from its knowledge. The key components of an expert system are an inference engine and a knowledge base. The separation of control (the inference engine) from knowledge (knowledge base) is a hallmark of an expert system. Other components of an expert system include a user interface, a knowledge-acquisition module, and an explanatory interface.

APPENDIX A

FUZZY LOGIC - A type of logic that attempts to deal with imprecise information. An extension of boolean logic. In boolean logic an entity is considered to be a member of a set or it is not. In fuzzy logic membership in a set is a matter of degree. A given object may belong to different classifications with varying degrees of certainty.

GAME PLAYING - A field of artificial intelligence that specializes in using games to implement and test artificial intelligence principles.

GENERAL PROBLEM SOLVING - The field of AI devoted to building programs which will solve a large class of problems. The programs in this category emphasized reasoning as opposed to knowledge. General Problem Solver is the best known example of this type of program. Such programs are in disfavor at this time because although they can solve a range of problems they are very different. They have been supplanted by expert system programs.

INTELLIGENT COMPUTER-AIDED INSTRUCTION - (ICAI) A program that attempts to teach a body of knowledge using a computer program. It has the following characteristics: It analyses a student's performance and tailors an instruction program for the student based on the analysis. It takes active control of the learning process. It generally consists of three components: a problem solving module, a student module, and a tutoring module. The problem-solving module contains the information that the student is to learn. The student module is a model of the program's perception of what the student knows. The tutoring module is responsible for choosing the teaching strategies to be followed.

KNOWLEDGE ACQUISITION - The procedure in AI of interacting with an external source, usually a domain expert, to find and organize knowledge for the purpose of transferring the knowledge to an expert system to solve problems. There are two types of knowledge acquisition--manual knowledge acquisition and automated knowledge acquisition. Manual knowledge acquisition refers to such procedures as a knowledge engineer interviewing a domain expert, and verbal protocol analysis. Manual knowledge acquisition can also be formally broken down into the stages of identification, conceptualization, formalization, implementation, and testing. Automated knowledge acquisition refers to the process of getting computers to learn for external sources, especially from experts.

KNOWLEDGE ENGINEER - A professional who interacts with a domain expert in order to obtain the necessary facts and relationships among the facts from the domain expert to build an expert system. He or she should know which tools to use and how to use them; have interpersonal skills; and be able to test expert systems, design interesting interfaces, construct well-organized knowledge bases, verify knowledge bases, and avoid combinatorial explosion.

KNOWLEDGE REPRESENTATION - A systematic means of organizing, portraying, and storing knowledge in a computer program which leads to knowledgeable behavior.

LOGIC PROGRAMMING - A type of programming based on logic that is relatively independent of the underlying structure of the machine on which it is operating. The heart of logic programming is making deductions from a set of facts. Logic programming consists of objects, how the objects are related, and operators.

MACHINE LEARNING - 1. Induction learning, a type of automated knowledge acquisition that makes the assumption that object descriptions are available and a computer program then uses examples concerning the objects to construct a decision tree which can than be used to produce rules. The process is based on statistical techniques and numerical algorithms. 2. A field of AI that attempts to build programs which learn from experience. This includes learning by induction, concept learning, learning in neural nets, discovery learning, learning by interaction, learning from instruction, learning by interaction, learning form instruction, learning by analogy, model-driven learning, skill refinement, and data-driven learning.

MACHINE TRANSLATION - The translation of text from one human language to another human language by computer. Results have been disappointing in this area.

MACHINE VISION - The reception, processing, and understanding of visual images in AI. The information processing chore of comprehending a scene from its images.

NATURAL LANGUAGE PROCESSING - A branch of artificial intelligence programming whose goal is to facilitate communications between humans and computers using written human language. It consists of two areas: natural language understanding and natural language generation. Three important areas of study in natural language processing are lexical analysis, syntactic analysis, and semantic analysis. Parsers are an important area of study in language understanding. Three common parsers are state-machine parsers, context-free parsers, and noise-disposal parsers. Speech synthesis and speech recognition are not considered a part of natural-language processing.

NEURAL NETWORKS - 1. A web of nerve cells in a living organism. 2. A computer simulation of the brain, which consists of at least one neuron and synapses. The neuron has an activation level and a transfer function. The synapses are the connection points for the neurons and are made up of an input, a weight, and an output. The neurons may be connected to one another in a complex net and they work in parallel with each other. 3. Self-organizing systems of simple interconnected processing units which possess a learning rule and are capable of learning.

PATTERN MATCHING - The search for similarities between symbolic expressions. The process of comparing a specific structure and a more general structure to see if the more specific structure is an instance of the more general structure. An example is matching a pattern against a data base. A pattern is a type of template that has variables.

PATTERN RECOGNITION - 1. The subfield of artificial intelligence that focuses on the recognition and classification of patterns. 2. Programs based on patterns of bits rather than reasoning. Example: Holland classifier. 3. The use of statistical techniques and templates to process and classify patterns of data. The process of classifying and image into a predefined category. The two stages of pattern recognition are feature extraction and pattern classification. The term should not be confused with pattern matching which is involved with symbolic manipulations rather than numeric computations.

PLANNING - Making preparations to take action. The ability of an expert system to organize a series of steps to reach a goal. The process of deriving a sequence of actions before proceeding with the action. The goal of planning is to reduce the search space as much as possible before problem solving begins. See also Hierarchical planning and repair; Least commitment; Nonhierarchical planner; Opportunistic planning.

ROBOTICS - The field of AI that pursues the goal of developing intelligent robots. Robots have a number of advantages: improved quality, greater productivity, ability to work in hazardous environments, and reduced costs.

SEARCH - A problem-solving procedure. The investigation of different potential solutions in the problem-solving or planning process. The procedure of exploring a search space in an attempt to solve a problem.

SPEECH RECOGNITION - The discrimination of speech sounds by a computer. Contrast with Speech-understanding system.

THEOREM PROVING - The subarea of AI involved in using computers to prove mathematical and logic theorems. Both deductive reasoning and intuitive leaps can be employed in theorem proving.

UNCERTAINTY - 1. The inability of an expert system to obtain a value. In that event the expert system may ask the user, attempt to calculate the value, use a default value, or attempt to continue reasoning without the value. 2. The lack of sure knowledge as to whether or not a rule, a fact, or a user response may be correct. Inexact reasoning is used to reason with uncertainty. Certainty factors, fuzzy logic, Shafer-Dempster theory, and Bayes' theorem are methods of inexact reasoning available for dealing with uncertainty.

UNDERSTANDING SYSTEMS - AI systems that are able to use a body of textual knowledge to answer questions about the body of knowledge. The more advanced AI understanding systems do more than simply find a piece of text in response to a query. Understanding systems are able to scrutinize text, make inferences, and answer questions when the answer is not explicit but only implied.

CRITICAL TECHNOLOGY ASSESSMENT OF THE U.S. ARTIFICIAL INTELLIGENCE SECTOR

APPENDIX B

AI ROOTS AND TRENDS

U.S. DEPARTMENT OF COMMERCE
BUREAU OF EXPORT ADMINISTRATION
OFFICE OF INDUSTRIAL RESOURCE ADMINISTRATION
STRATEGIC ANALYSIS DIVISION

APPENDIX B

AI ROOTS AND TRENDS

AI technology arose in the United States and Great Britain after World War II as an outgrowth of computer science. Most subsequent AI research and development took place in the United States, although important work was also done at the University of Edinburgh in Scotland, the University of Marseilles in France, and several other foreign universities. In the United States, AI research proceeded as part of a grander "Cold War" strategy to

HISTORY OF ARTIFICIAL INTELLIGENCE AT A GLANCE	
1.	People have been drawn to the idea of artificial intelligence since ancient times, centuries before the term was coined. The advent and advances in computer technology provided the experimental "vehicle" to seriously examine the "intelligence methodologies" put forth in this field.
2.	In 1950, British mathematician Alan Turing, the "Father of Artificial Intelligence," suggested an imitation game, known as the Turing Test, as a means of ascertaining if a machine can "think." Turing built the first AI programs during World War II to break German codes.
3.	The AI revolution was "officially" launched in 1956 at the Dartmouth Conference, which was organized by John McCarthy of Dartmouth, Marvin Minsky of Harvard, Nathaniel Rochester of IBM, and Claude Shannon of Bell Labs.
4.	Other notable participants in the Dartmouth Conference included Allen Newell and Herbert Simon of the Carnegie Institute of Technology and RAND Corporation, Arthur Samuel and Alex Bernstein of IBM, and Oliver Selfridge of MIT.
5.	Although AI research is conducted at many colleges and universities, the primary academic centers of AI research are the three schools that were early leaders in the field: Carnegie-Mellon University (CMU), the Massachusetts Institute of Technology (MIT), and Stanford University.
6.	Many early AI developments resulted from research into games, such as checkers and chess.
7.	The U.S. Government, particularly the Defense Department's Advanced Research Projects Agency, has been the major sponsor of AI research.
8.	Many new companies were formed during the 1980s to commercialize AI technology.
9.	Neural nets developed to a practical level by parallel distributed processing group at University of California, San Diego during 1980s.
10.	The commercial potential of AI is being examined and put to use by major corporations, including AT&T, Digital Equipment, DuPont, General Electric, Ford Motor, McDonnell Douglas, Intel, Schlumberger, Texas Instruments, IBM, and Xerox.

Source: *Understanding Artificial Intelligence*, 1988, developed from Chapter 2: History of Artificial Intelligence.

maintain a technological lead over the Soviet Union. A large portion of AI research funding came from the Department of Defense, particularly the Advanced Research Projects Agency (ARPA). ARPA was created in 1958, as a direct consequence of the Soviet Union's launch of Sputnik in October, 1957. In 1993 ARPA funded about $70-80 million of AI research, over half of the U.S. Government total.

APPENDIX B

B.1 Alan Turing and the Turing Test

Alan Turing was a British mathematician best remembered for his theoretical computing devices that bear his name (Turing machines). While a graduate student at Princeton University in 1936, Turing published "On Computable Numbers," a paper in which he conceived of a machine that could move from one state to another by following a rigorous set of rules. This led to a computing scheme that foreshadowed the logic of digital computers. During World War II, Turing worked in the British Foreign Office, where he played a leading role in breaking enemy codes. He later worked on the development of an electronic computer and the application of mathematical theory to biological forms, as well as developing his theories of artificial intelligence.

In 1950, Turing wrote a provocative article entitled "Computing Machinery and Intelligence," which secured for him the distinction of being generally recognized as the "father" of artificial intelligence. Turing began the article with a proposition, "I propose to consider the question, can machines think?" Turing foresaw that there would be many objections to the proposition that machinery could produce thought. Recognizing that semantic difficulties alone could render it impossible ever to answer the question satisfactorily, Turing suggested a test, in the form of a game, that could help to decide the issue. He called it the "imitation game." Later, in his honor, it became known as the "Turing Test."

Simply stated, the Turing test suggests that a prerequisite for artificial intelligence should be that the computer be indistinguishable from a human in its thought patterns. An example would be if playing two games of chess through the mail, one with a human chess master and one with a computer, one should not be able to distinguish which is which.

The general concept of a Turing Test has evolved with the passage of time. Today, a Turing Test is considered to be any situation in which a human converses with an unseen respondent and attempts to determine if the dialogue is being conducted with a human or a computer. If a computer can fool you into believing that you are talking to a human, one school of thought holds that the computer can be said to be intelligent.

The idea of such a test stimulated research in the area. It also revealed the true difficulty of putting "common sense" into a machine. One well-known example of this test is a program called ELIZA, developed by Joseph Weizenbaum at the Massachusetts Institute of Technology (MIT). ELIZA simulated a psychiatrist. For example, if you were to tell ELIZA that you missed your children, ELIZA might respond with, "Why do you miss your children?" or "Tell me more about your family." Either response might lead you to believe that ELIZA understands what you are saying, when actually it was using some programming tricks to construct responses from your statements. Many people were fooled; however,

APPENDIX B

ELIZA did not in actuality pass the Turing Test.

While some no longer consider the Turing test a valid indicator of machine intelligence, its spirit endures. Many people in AI continue to maintain that when we can simulate intelligent behavior so closely that it is <u>impossible</u> for even the most discerning individuals to tell the difference between a human and a computer, it will be fair to say that we have created a machine that thinks.

B.2 The Dartmouth Conference

In the summer of 1956, the small college town of Hanover, New Hampshire hosted the first AI conference at Dartmouth College. The Dartmouth Conference lasted about two months with ten people participating at various times. The disciplines represented included mathematics, neurology, psychology, and electrical engineering. The group was linked by the computer, which they were using to search through their disciplines for ways to simulate human intelligence. There was no universal agreement about what to call the new science; however, artificial intelligence, the name suggested by John McCarthy, one of the conference organizers, has come to be associated firmly and irrevocably with the field.

The conference, funded by a $7500 Rockefeller Foundation grant, was organized by four scientists, two from academia and two from industry. According to their proposal, the conference intended to explore the possibility

> **"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."**

That possibility remains the central focus of AI research. Thus began the quest for artificial intelligence.

B.4 AI's Pioneers

While the Dartmouth Conference achieved little other than naming the field and defining its goal, it is especially noteworthy for the distinguished people it brought together. The major pioneers of AI research in the United States were all participants in the conference. Even today, the leadership of the American AI community is composed largely of conference participants, their students, and their students' students.

The four scientists who organized the Dartmouth Conference were:

1) John McCarthy, an assistant mathematics professor at Dartmouth. McCarthy was, in fact, the primary moving force behind the conference.

2) Marvin Minsky, a junior fellow in mathematics and neurology at Harvard University. Minsky had been a graduate student with McCarthy at Princeton University; both had subsequently worked briefly with Claude Shannon at Bell Labs.

3) Nathaniel Rochester, Manager of Information Research for IBM. Rochester had a keen interest in intelligent machines, especially in getting machines to exhibit original behavior in problem solving.

4) Claude Shannon, of Bell Labs, had established his reputation firmly in the field of information science.

McCarthy and Minsky continue to be leaders of the AI community. Rochester and Shannon, however, did not devote themselves to AI after the Dartmouth Conference.

John McCarthy is currently a professor of computer science at Stanford University. He was responsible for several major developments in artificial intelligence, and in computer science in general. For example, in 1958 he invented LISP,[1] until recently the most commonly used AI programming language in the United States. He also developed the concept of "time sharing" to better utilize the smaller and much more expensive computing capacities of earlier computers.

With John McCarthy, Marvin Minsky cofounded the Artificial Intelligence Laboratory at MIT, one of the most prestigious centers of artificial intelligence research in the world. Minsky served as the AI lab's director for a few years. He is also a past president of the American Association for Artificial Intelligence (AAAI).

Minsky is perhaps best known for his work in the organization and representation of

[1]LISP, an acronym for List Processing Language, is a high-level computer language used widely in the study and application of artificial intelligence. LISP became popular with researchers at the Massachusetts Institute of Technology, who exploited its logical and economic structure in their studies of AI. LISP is an interactive language in which the user builds complex commands, or "lists," by putting together a series of related words, or "atoms," by means of expanding parentheses.

knowledge structures, an area with important implications throughout the various AI technologies. He has been a prolific writer in both technical and general publications, and remains one of the more outspoken proponents of AI. Minsky made a devoted study of the inner workings of the mind, while many people continued to believe that intelligence was something conscious and physical that controlled our actions. In a recent work, Minsky suggests a new theory of the mind that makes what he calls the "meat machine" (as opposed to the electronic machine) more plausible than ever. In *The Society of Mind*, Minsky describes the mind as a society of "agents" that each perform simple tasks. When these agents work in concert, they give the perception of intelligent action to living creatures. As these agents learn small reactions to environmental changes, the whole being exhibits the skill of learning and the art of self-awareness.

In addition to the Dartmouth Conference organizers, others who attended included: Allen Newell and Herbert Simon, sometimes referred to as the Carnegie-RAND duo; Arthur Samuel and Alex Bernstein of IBM; and Oliver Selfridge of MIT. Newell and Simon were both on the faculty of the Carnegie Institute of Technology (now Carnegie-Mellon University or CMU) and on the staff of the RAND Corporation. Independently, Samuel and Bernstein were programming computers to play games, one of the first areas of AI research. Selfridge had served as Norbert Wiener's assistant during the preparation of *Cybernetics*.[2] Selfridge's presentation at RAND Corporation of a pattern-recognition program first interested Allen Newell in the AI field.

Newell and Simon, along with McCarthy and Minsky, are now considered to be the leading American AI pioneers. Herbert Simon is an American social scientist and computer theorist who was awarded the 1978 Nobel Prize for economics for his "research into the decision-making process within economic organizations." He first studied management. Foreseeing the role computers were to occupy in business, he wrote *The New Science of*

---

[2]Cybernetics is the title Norbert Wiener gave his 1948 book on the subject. Wiener developed a new approach to understanding the workings of the universe. Wiener suggested a model that has proven extremely valuable in understanding computers as well as people - he suggested the transfer of "information" rather than energy (Newton's universe), as a way to model many different kinds of scientific phenomena. Cybernetics (Ancient Greek for helmsman or controller) was the name Wiener used to describe his informational approach. By describing interrelated systems in terms of exchanging information, cybernetics points out the functional similarities between humans and machines, similarities which troubled Wiener, but that proved to be invaluable in early artificial intelligence research.

Management Decision (1960), which analyzes the uses of computers in management.[3]

In a series of papers beginning in 1957, written in collaboration with Allen Newell and J. C. Shaw, Simon developed heuristic (rule-of-thumb) problem-solving programs that employ strategies that are not mere mechanical search procedures. The achievements of these programs were influential in the development of cognitive psychology and artificial intelligence in computer research, fields in which his studies of human problem solving have become the dominant model.

Allen Newell, the junior partner of the Newell and Simon team, left the graduate program in mathematics at Princeton in order to join the RAND staff. His work at RAND brought him in contact with logistical organizations in the Department of Defense where he became interested in studying the way people behaved in organizations. It occurred to Newell that the study of organizational behavior could be approached scientifically by modeling human behavior on a computer, and he became interested in the relationship between the computer and the brain. Like Marvin Minsky, Allen Newell has served as president of the AAAI (in fact, Newell was the first president of the organization).[4]

Allen Newell and Herbert Simon came to the Dartmouth Conference with something that no one else at the conference had - a working AI program they called the *Logic Theorist*. In collaboration with J. C. Shaw, a RAND computer scientist, Newell and Simon had developed the Logic Theorist to generate proofs of mathematical theorems. Specifically, the theorems that the Logic Theorist was designed to prove were those of *Principia Mathematica*, the early 20th-century proposition calculus (logical inference) system of Bertrand Russell and Alfred North Whitehead.

In addition to the Logic Theorist, Newell and Simon also collaborated with J. C. Shaw on the development of IPL (Information Processing Language), the first symbolic programming language and a precursor to Prolog and, to a lesser degree, LISP. They also created the *General Problem Solver*, an early AI program that solved problems by choosing an appropriate path to a specific goal ("means-ends analysis").

---

[3]Other books by Simon include *Administrative Behavior* (1947); *The Sciences of the Artificial* (1969, 2d ed. 1981); with Allen Newell, *Human Problem Solving* (1972); *Models of Bounded Rationality and Other Topics in Economics* (1982); *Reason in Human Affairs* (1983), and *Models of My Life* (1991).

[4]Allen Newell coauthored *Human Problem Solving* (1972), and authored *The Psychology of Human-Computer Interaction* (1983), and *Unified Theories of Cognition* (1990).

APPENDIX B

In the late 1950s and early 1960s, scientists Newell, Simon, and J. C. Shaw created their "logical theorist" computer program, and introduced symbolic processing.

Instead of building systems based on numbers, they attempted to build systems that manipulated symbols.

Their approach was powerful and is fundamental to most work in AI to this day. In this approach, knowledge is expressed as rules, for example, "If x is a bird, then x can fly." If such an AI system determines or is told that a robin is a bird, then it will infer that the robin can fly.

B.5 Early Research

In 1937, Claude Shannon, a graduate student at MIT, used Boolean algebra[5] to describe the operation of electrical switching circuits.

Shannon's ideas contributed to the developing field of information science and led directly to the binary system of information storage used in the digital computer.

While working at Bell Labs in 1953, Shannon wrote, "The problem of how the brain works and how machines may be designed to simulate its activity, is surely one of the most important and difficult questions facing current science." He suspected that if Boolean algebra, a representation of human thought, could be used to describe electrical circuits, then

---

[5]The English mathematician George Boole (1815-1864) is best known for his work in mathematical logic. Much of his work on logic, algebra, and probability is contained in *An Investigation of the Laws of Thought on Which Are Founded the Mathematical Theories of Logic and Probabilities* (1854). Boole approached logic in a new way: he reduced it to a simple algebra, thereby incorporating logic into mathematics and giving the development of logic a new direction. Boole's two-valued, or binary, algebra is the simplest form of the more general Boolean Algebra, which he developed, and has wide applications in the design of modern computers.

perhaps circuits could be used to describe thought. The computer's role as an "experimental vehicle" for testing and evaluating these ideas was central to the field.

Scores of AI systems have been built as a means for uncovering and facing the problems of producing intelligent behavior by a computer.

Much early AI research was absorbed in game theory. In attempting to create games, such as checkers and chess, that could successfully compete against humans, researchers had to consider how a computer might reason out the best of a number of possible moves. Realizing that the human chess player does not mentally run through all combinations of movement, developers gave consideration to the methods by which a chess master chooses a move. Incorporating notions of extrapolation and deduction, computer game theory has developed so that computers can regularly out-play human players.

In 1947, one year after leaving Bell Labs, Arthur Samuel proposed a checker-playing machine. The purpose of the machine was to generate publicity in an effort to raise funds to build a computer at the University of Illinois. He assumed that it would be a trivial task to build a small machine that could play championship-level checkers; he hardly suspected that he was embarking on a 20-year project.

When Samuel left Illinois for IBM, he continued to think about a checker playing machine. As he became involved in the design of IBM's first general-purpose computers, he began to write, and continually rewrite, the checkers program he had conceptualized at Illinois. The feature of Samuel's Checker Player that has earned him a place in AI history is that it actually learned from its mistakes, and it may have been the first program ever to do so. By 1961, the program played at the level of a checkers master; and it could beat Samuel, who ironically was never very good at checkers.

Samuel had disproved the widespread belief that a computer can only do exactly what it is programmed to do and thus can never exhibit original behavior.

The game of chess has also received a great deal of attention. Claude Shannon wrote an article that appeared in *Scientific American* in 1950, in which he discussed the possibility of using computers to play chess. Shannon may have been the first to point out that having a computer consider every possible combination of moves was not a practical chess-playing strategy. He estimated that even if it could evaluate one million moves per second, a computer would take 10^{95} (i.e., 10 to-the-95th-power) years to select a move.

This is known as the "Combinatorial Explosion," the exponential increase in the number of ways the objects in a set can combine as the number of objects increases. This is a major problem in artificial intelligence because the rapid increase in different combinations causes such a large increase in the search space that it is not practical to be able to search each alternative within a reasonable period of time.

Alex Bernstein, an enthusiast of chess, wrote a chess playing program. While working at IBM, Bernstein used sophisticated heuristics to improve the level of play of his program. By the time the program was complete, it was able to play at the novice level. His chess program is important in AI history because of the heuristic techniques it employed to search for the best moves.

The complexity of chess required Bernstein to explore, in great detail, various methods of eliminating possibilities, a concept central to AI research.

In 1957, a chess program was written at MIT. It worked by selecting the seven best moves, the seven best replies, and the seven best replies to each of these replies, creating a tree of 2,401 (i.e., 7^4) variations. In 1967, a match was played between a U.S. and Soviet program. Since then tournaments have been conducted regularly between competing programs, which are becoming a progressively greater challenge to human players as well. In 1983, a program named Belle, designed at AT&T's Bell Laboratories, became the first to reach the U.S. master level of ability, and in 1988 an IBM-designed program, Deep Thought, defeated one grand master and tied another. Only the international grand master level remains to be toppled, and programs have already won some games at that level in exhibitions where several challengers were taken on simultaneously.

B.6 Development of Expert Systems

One of the most useful ideas that emerged from AI research is that facts and rules, *declarative knowledge (nouns)*, can be represented separately from decision-making algorithms, *procedural knowledge (verbs)*. This realization has had a profound effect both on the way that scientists approach problems and on the engineering techniques used to produce AI systems.

By adopting a particular procedural element, or inference engine, development of an AI system is reduced to obtaining and codifying sufficient rules and facts from the problem domain.

APPENDIX B

This codification process is called *knowledge engineering*.

In the mid-1960s, the first knowledge-based expert program was written. Called *DENDRAL*, it could predict the structures of unknown chemical compounds based on routine spectographic analysis. DENDRAL was developed by Ed Feigenbaum, the inventor of expert systems, and others at Stanford. More sophisticated rule-based expert systems were subsequently developed, notably the *MYCIN* program. MYCIN uses rules derived from the medical domain to reason backward (deduce) from a list of possible diseases that fit the symptoms of the current case. Many expert systems of similar design have been constructed.

Many of the early expert systems were developed in the field of medicine.[6] A 1986 survey conducted by Waterman and reported in his book *A Guide to Expert Systems*, showed 30 percent of the applications medically related. By 1993, after explosive growth in other areas, the medical field accounted for about 12 percent of the applications. The fastest growth occurred in business and industrial areas, where from 1986-1993, the percent of applications rose from 10 to 60 percent. The systems are now found in virtually all industries, and nearly all areas where human expertise is spread thinly.[7]

About 30 percent of expert systems are used for diagnostic purposes, although this ratio may be declining as other areas continue developing. Diagnostics were an early and obvious use of expert systems, in part because of the role human experts play in the area. They are also easier to develop because most diagnostic problems have a finite number of solutions and require a limited amount of information to reach a solution. One other explanation may be the low risk preference of organizations faced with new technologies. Because they are proven and easy to build, they make attractive candidates for firms venturing into the field.[8]

Reducing expert system development to knowledge engineering has opened the door to non-AI practitioners. In addition, business and industry have been recruiting AI scientists to build expert systems. However, an impediment to building even more useful systems is the problem of input, in particular, the feeding of raw data into an AI system. To this end, much effort is currently being devoted to speech recognition, character recognition, machine vision, natural-language processing, and better interfaces between expert systems and databases. A second problem is in obtaining knowledge. It has proved arduous to extract knowledge from an expert and then code it for use by the machine. To this end, efforts are currently being devoted to learning and knowledge acquisition.

---

[6]Liability and implementation issues limited their actual clinical deployment.

[7]*Expert Systems Catalog of Applications*, 1993, by John Durkin, p. v.

[8]Ibid. p. viii.

CRITICAL TECHNOLOGY ASSESSMENT OF THE U.S. ARTIFICIAL INTELLIGENCE SECTOR

APPENDIX C

SURVEY INSTRUMENT

U.S. DEPARTMENT OF COMMERCE
BUREAU OF EXPORT ADMINISTRATION
OFFICE OF INDUSTRIAL RESOURCE ADMINISTRATION
STRATEGIC ANALYSIS DIVISION

APPENDIX C

Ref. # 65

U.S. Department of Commerce
Bureau of Export Administration

OMB Control # 0694-0069
Expires:12/31/92

CRITICAL TECHNOLOGY ASSESSMENT: ARTIFICIAL INTELLIGENCE

THIS REPORT IS REQUIRED BY LAW

Failure to report can result in a maximum fine of $1,000 or imprisonment up to one year, or both. Information furnished herewith is deemed confidential and will not be published or disclosed except in accordance with Section 705 of the Defense Production Act of 1950, as amended (50 U.S.C. App. Sec. 2155).

BURDEN ESTIMATE AND REQUEST FOR COMMENT

Public reporting burden for this collection of information is estimated to average 4.5-10 hours per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to BXA Reports Clearance Officer, Room 4513, Bureau of Export Administration, U.S. Department of Commerce, Washington, D.C. 20230, and to the Office of Management and Budget, Paperwork Reduction Project (0694-XXXX), Washington, D.C. 20503.

GENERAL INSTRUCTIONS

1. Please complete this questionnaire as it applies to your organization or firm. The reverse side of this page contains a table that designates which questions you are to answer, depending on your firm size, or whether you are an academic institution or government agency. Please review this table carefully. If you are not involved in artificial intelligence, check the bottom line of the table, then sign the certification on last page and return this questionnaire to the address below. Otherwise, **Your response is due by December 31, 1992.** The survey has six parts as follows:

   PART I Identification
   PART II Competitive Considerations
   PART III Commercial Information
   PART IV Employment Information
   PART V Research and Development
   PART VI Technology

2. It is not our desire to impose an unreasonable burden on any respondent. IF INFORMATION IS NOT READILY AVAILABLE FROM YOUR RECORDS IN THE FORM REQUESTED, FURNISH ESTIMATES AND DESIGNATE BY THE LETTER "E". Report calendar year data, unless otherwise specified in a particular question. Please make photocopies of forms if additional copies are needed.

3. Questions related to this questionnaire should be directed to Mr. John Tucker, Senior Industry Analyst, (202) 377-3984, U.S. Department of Commerce.

4. Before returning your completed questionnaire, be sure to sign certification on the last page and identify the person and phone # to contact should we have follow-up questions. Return completed questionnaire by December 31, 1992 to:

   U.S. Department of Commerce
   BXA/OIRA, Attn: Brad Botwin, Director,
   Strategic Analysis Division, Room H3878
   Washington, D.C. 20230

APPENDIX C

ARTIFICIAL INTELLIGENCE SURVEY CONTENTS

APPENDIX C

CODES FOR MAJOR CATEGORIES OF ARTIFICIAL INTELLIGENCE

These Codes should be used in completing the following questions:

PART III - 1. SHIPMENTS (page 9)
and
PART V - 2. AI RESEARCH PROJECTS (page 16)

CODES FOR MAJOR CATEGORIES AND SUBGROUPS OF ARTIFICIAL INTELLIGENCE	
LETTER CODES FOR MAJOR CATEGORIES	SAMPLES OF SUBGROUPS WITHIN MAJOR CATEGORIES
A. APPLICATIONS	1. AUTOMATED PROGRAMMING 2. EXPERT DATA BASE SYSTEMS 3. EXPERT SYSTEMS 4. GAME PLAYING
B. AUTOMATED REASONING	1. DISTRIBUTED PROBLEM SOLVING 2. FUZZY LOGIC 3. GENERAL PROBLEM SOLVING 4. PATTERN MATCHING 5. REASONING 6. SEARCH 7. UNCERTAINTY
C. COGNITIVE MODELING	1. DECISION-MODELING SYSTEM 2. UNDERSTANDING SYSTEMS
D. KNOWLEDGE REPRESENTATION	1. KNOWLEDGE REPRESENTATION 2. NEURAL NETWORKS
E. LEARNING	1. KNOWLEDGE ACQUISITION 2. MACHINE LEARNING 3. PATTERN RECOGNITION
F. NATURAL LANGUAGE UNDERSTANDING	1. MACHINE TRANSLATION 2. NATURAL LANGUAGE PROCESSING 3. SPEECH RECOGNITION
G. PLANNING AND PROBLEM SOLVING	1. PLANNING
H. ROBOTICS AND MACHINE VISION	1. ROBOTICS 2. MACHINE VISION

APPENDIX C

PART I IDENTIFICATION

1. **NAME AND LOCATION**: Enter your organization's name, the area or department in your organization involved in artificial intelligence (AI), and your address below.

Name:____

AI Area:______________________________________

Address:______________________________________

City:____

State:_________________________Zip Code:________

2. **OWNERSHIP**: If your organization is a business firm that is wholly or partly owned by another firm, indicate the name and address of the parent firm, extent of ownership, and when you were acquired below.

Name:____

Address:______________________________________

City:____

State:_________________________Zip Code:________

Country (if not U.S.):____________________________

Percent Ownership:________% When Acquired (year)?_____

3. **ESTABLISHMENTS**: On the table below, please identify the location of each establishment in which your organization conducts AI activities. Also, indicate the primary activity of each establishment in the last column as: a) engaged in R&D, b) AI development/production, c) consulting, d) other (specify). (See definition of establishment.)

Establishment Locality	State	Zip Code	Primary Acitivity

PART I (continued)

4. **MERGERS AND ACQUISITIONS:** Please describe any mergers or acquisitions your organization has been involved in since January 1 1989, with respect to your AI activities.

____

____

____

____

____

5. **COOPERATIVE AGREEMENTS:** Please complete the table below, listing the consortiums, joint ventures, strategic partnerships or alliances, teaming efforts, licenses, marketing agreements, or other arrangements with other organizations you have entered into with respect to your AI activity. List the agreements with U.S. entities in the top portion of the table, and those with foreign entities on the bottom portion.

Agreements with U.S. Organizations	U.S. Partner's Name	Primary AI Activity

Agreements with Foreign Organizations	Foreign Partner's Name	Country	Primary AI Activity

APPENDIX C

PART II COMPETITIVE CONSIDERATIONS

1. **IMPROVED COMPETITIVENESS**: How has AI improved the competitiveness of your company, or one of your clients? (If available, please provide before and after data that shows cost savings, productivity increases, or the ability to do things not feasible without AI.)

____

____

____

____

____

____

2. **APPLICATIONS**: a.) What are three major **defense applications** of AI your organization created or helped create, and what do they do?

____

____

____

____

____

____

b.) What are three major **commercial applications** of AI your organization created or helped create, and what do they do?

____

____

____

PART II (continued)

3. **DEFENSE CONVERSION AND DUAL USE:** a.) Do AI programs made for defense applications, also have applications in the commercial sector? yes _____, no_____

Please comment:____

____

Are you aware of any federal, state or local government programs to assist firms in converting defense-related operations to commercial operations? What kinds of programs would be useful?

____

____

____

b.) To your knowledge, has the Department of Defense made an effort to promote dual uses of AI?

yes _____, no _____

If yes, please explain how:____

____

____

____

4. **IMPACT OF DEFENSE CUTS:** What impact has (or will) the reduction or termination of major defense programs had (have) on your organization?

____

____

____

____

____

APPENDIX C

PART II (continued)

5. **GOVERNMENT ROLE**: How could the Federal or State Governments assist the AI sector? (e.g., change tax laws, change legal environment, change export controls, government procurement, fund research, promote cooperative agreements, etc.)

____

____

____

6. **WHERE IS UNITED STATES LOSING GROUND**: If you believe the United States is falling behind it international competitors in any area of AI research, production, or applications, please identify the area, who is leading (or will be shortly), and why we are falling behind.

____

____

____

7. **MARKET CONSTRAINTS**: What factors would improve the demand for your product?

____

____

____

What problems have you encountered in selling AI to your customers?

____

____

____

What factors would make AI more attractive to your customers?

____

____

PART II (continued)

8. **STRATEGIC PARTNERSHIPS:**

A. Do you perceive the U.S. antitrust laws to be a barrier to strategic alliances:

1. With other U.S. firms yes______, no______
2. With foreign firms yes______, no______
3. In horizontal relationships yes______, no______
4. In vertical relationships yes______, no______

B. Have you had actual experience in which U.S. antitrust laws have created a barrier to cooperation with other firms:

1. In R&D partnerships yes______, no______
2. In manufacturing relationships yes______, no______

If yes, please describe the experience:____

____

____

C. Do you currently have or have you in the past had vertical alliances with suppliers, manufacturers, or distributor firms in your field? Please elaborate whether these involve: 1. Foreign firms, 2. R&D, 3. Manufacturing, 4. Marketing, 5. Short (1-5 years) or Long term (5 years or more)

____

____

____

____

D. Would you consider such alliances in the future? yes______, no______

If not, why not?____

____

APPENDIX C

PART III COMMERCIAL INFORMATION

1. **SHIPMENTS/RECEIPTS** - Please complete the following table by reporting 1989-1991 dollar values (in $000s) of shipments (receipts) for domestically developed AI products or AI services: a) sold by your firm to unrelated parties (enter on upper table), and/or b) market value equivalent of AI developed by your firm for in-house use (on lower table). Report separately for each major category of AI shown by letter code on page ii applicable to your operations. Please **ENTER** the letter of the category being reported where requested at the top of each table. **(If you have receipts in more than one category, please make copies of this page for each additional category you ship.)**

a) **SHIPMENTS OF ARTIFICIAL INTELLIGENCE SOLD TO UNRELATED PARTIES:**

Letter Code of Category Being Reported:_______(See Codes, page ii)

AI SOLD TO:	1989	1990	1991
Department of Defense			
Other Government			
Manufacturers			
Other Business			
Total			
Exports included in above Total			
to Japan			
to Western Europe			
to Other			

b) **IN-HOUSE DEVELOPMENT OF ARTIFICIAL INTELLIGENCE FOR IN-HOUSE USE**

Letter Code of Category Being Reported:_______(See Codes, page ii)

AI DEVELOPED FOR:	1989	1990	1991
Defense Applications			
Manufacturing Operations			
Other Operations			

APPENDIX C

PART III (continued)

SPECIAL INSTRUCTION: If AI revenues are less than 50 percent of total revenues for your firm or division, check here_____, and respond only to b.) for questions 2 and 3 below.

2. **PROFITABILITY**: Please enter the financial information (in $000s) as specified below for 1989-1991 for a.) your parent firm (the corporate level, if applicable), and b.) the dollar amounts that apply exclusively to your AI production operations in the United States. **Please photocopy this page if both a.) and b.) apply.**

	1989	1990	1991
Net Sales (1)			
Cost of Goods Sold (2)			
Operating Income (3)			
Net Income before taxes (4)			
Aftermarket Revenues			

(1)Sales and service fees
(2)Includes material purchases, direct labor, and other lab costs such as depreciation and inventory carrying costs.
(3)Difference between Net Sales and Cost of Goods Sold
(4)Operating income less general, selling and administrative expenses, interest expenses and other expenses, plus other income

3. **FINANCIAL BALANCES**: Please provide end of year balance sheet information (in $000s) as specified below for 1989-1991 for a.) your parent firm (the corporate level, if applicable), and b.) the dollar amounts that apply exclusively to your AI production operations in the United States. **Please photocopy this page if both a.) and b.) apply.**

	1989	1990	1991
Current Assets			
Current Liabilities			
Inventories			
Total Assets			
Short Term Debt (1)			
Long Term Debt (2)			

(1) Principal payable in less than one year
(2) Principal payable in more than one year

PART III (continued)

4. **INVESTMENTS**: a.) Enter AI related expenditures for laboratories, software, hardware and equipment (in $000s) from 1989-1991. b.) If any of your investment expenditures were funded by outside sources such as the government, major firms, universities or others, please identify the amount and source below the shaded areas.

Investments in:	1989	1990	1991
Laboratories			
Hardware/Equipment			
Software			
Other*			
Total			
Outside Funders:			
1.__________________			
2.__________________			

* Specify:________________________

APPENDIX C

PART IV EMPLOYMENT

1. **EMPLOYMENT**: Please complete the table below by entering the number of employees by occupation category at your organization that are involved in AI research or production for (end-of-year) 1989-1991. If you have an occupation title commonly distinguished from the others, you may specify it where shown in the left column.

Occupation Category	1989	1990	1991
Knowledge Engineers			
Software Developers or Programmers			
Research Scientists			
Other: Specify__________________			
Other (clerical, administrative, support...)			
Total			
Number of Phds in above Total			
Number of trainees in above Total			
Training expenditures (in $000s)			
Total Payroll (salaries+wages+bonuses)			

2. **PROFESSIONAL DISCIPLINES**: Please complete the table below by entering the number of employees by professional discipline included in the occupational categories reported above.

Discipline	1989	1990	1991
Computer Scientists			
Mathematicians			
Cognitive Scientists			
Engineers			
Other (specify:__________________)			
Other (specify:__________________)			

APPENDIX C

PART IV (continued)

3. **TRAINING TIME**: Based on your experience, how long in terms of months (min/max) does it take for a college degree holder (or equivalent) with no previous AI experience to become proficient in the following occupation categories:

	Minimum		Maximum	
Knowledge Engineer -	_______	to	_______	months
Software Developer -	_______	to	_______	months
Research Scientist -	_______	to	_______	months

4. **LABOR SHORTAGES**: In the last five years have you experienced any shortages in the availability of or hiring qualified people? yes_____, no_____

If yes, please describe:

APPENDIX C

PART V RESEARCH AND DEVELOPMENT

1. **RESEARCH AND DEVELOPMENT**: Please enter research and development expenditures from 1989-1991. Enter separately the dollar amounts (in $000s) financed by your organization (in-house), a federal, state or local government agency, a business, thru a joint venture, or other entity. Also, for each category other than in-house enter any dollar amounts funded by a foreign source.

RESEARCH AND DEVELOPMENT - Research and development includes basic and applied research in the sciences and in engineering, and design and development of prototype products and processes. For the purposes of this questionnaire, research and development includes activities carried on by persons trained, either formally or by experience, in the physical sciences including related engineering, if the purpose of such activity is to do one or more of the following things:

1. Pursue a planned search for new knowledge, whether or not the search has reference to a specific application.
2. Apply existing knowledge to problems involved in the creation of a new product or process, including work required to evaluate possible uses.
3. Apply existing knowledge to problems involved in the improvement of a present product or process.

Source of Funding	1989	1990	1991
In-House			
Government: Federal a. Dept. of Defense	a._______	a._______	a._______
b. Other Federal	b._______	b._______	b._______
c. State & Local	c._______	c._______	c._______
d. Foreign Gov't	d._______	d._______	d._______
Manufacturing Firms: a. Domestic	a._______	a._______	a._______
b. Foreign	b._______	b._______	b._______
Non-Manufacturing Firms: a. Domestic	a._______	a._______	a._______
b. Foreign	b._______	b._______	b._______
Joint Ventures: a. Domestic	a._______	a._______	a._______
b. Foreign	b._______	b._______	b._______
Other:name__________ a. Domestic	a._______	a._______	a._______
b. Foreign	b._______	b._______	b._______
Totals:			

PART V (continued)

2. **AI RESEARCH PROJECTS**: Please complete the table below for each AI research and development project your organization is currently involved in, by answering the questions posed in the left hand column in the space provided in the right column. You may use the letter codes to identify the AI area being researched. **(Please photocopy this page for each project.)**

What is Project name?			
Whom is the Project for?			
What are starting and ending dates?	Start Date:	End Date:	
What is source and amount of funding?	Funding Source:	Funding Amount:	
What percent of research is basic, applied, etc.?	Percent Basic: ______%	Percent Applied: ______%	Percent Other: ______%
How many people are assigned to project?			
What area(s) of AI is being researched? (You may use Codes, page ii)			
What new (if any) developments are targeted?			

What are main objectives of project?

____

____

____

____

APPENDIX C

PART VI TECHNOLOGY

1. **SOFTWARE DEVELOPMENTS:** a.) What software developments have occured in the last five years that have enhanced your ability to produce AI, and/or conduct AI research?

b.) What software developments do you expect by the year 2000 that will further enhance your ability to produce AI, and/or conduct AI research?

2. **HARDWARE DEVELOPMENTS:** a.) What hardware and/or equipment developments have occured in the last five years that have enhanced your ability to produce AI, and/or conduct AI research?

b.) What hardware and/or equipment developments do you expect by the year 2000 that will further enhance your ability to produce AI and/or conduct AI research?

APPENDIX C

CERTIFICATION

The undersigned certifies that the information herein supplied in response to this questionnaire is complete and correct to the best of his/her knowledge. The U.S. Code, Title 18 (crimes and Criminal Procedure), Section 1001, makes it a criminal offense to willfully make a false statement or representation to any department or agency of the United States as to any matter within its jurisdiction.

____________ (Date)	______________________________ (Signature of Authorized Official)
____________ AreaCode/Telephone Number)	______________________________ (Type or Print Name and Title of Authorized Official)
____________ (AreaCode/Telephone Number)	______________________________ (Type or Print Name and Title of Person to Contact re this report)

GENERAL COMMENTS

Please use the space below to provide any additional comments or information you may wish regarding your operations, or other related issues that impact your organization.

APPENDIX D

RESEARCH STATISTICS

U.S. DEPARTMENT OF COMMERCE
BUREAU OF EXPORT ADMINISTRATION
OFFICE OF INDUSTRIAL RESOURCE ADMINISTRATION
STRATEGIC ANALYSIS DIVISION

APPENDIX D

RESEARCH STATISTICS

TABLE 1. ARTIFICIAL INTELLIGENCE
AS REPORTED RESEARCH PERFORMED BY COMPANIES BY SOURCE OF FUNDING, 1989-1994
(in thousands of dollars)

FUNDING SOURCE	1989	1990	1991	1992	1993	1994
Internal	$30,147	$27,640	$36,076	$30,329	$33,251	$37,924
Dept. of Defense	18,794	24,586	23,566	25,492	25,839	27,699
Other Federal Govt	3,900	3,320	4,799	4,000	4,200	4,200
State and Local Govt	0	0	0	450	0	0
Foreign Govt	500	500	600	700	0	0
Manufacturers-US	3,948	5,066	6,344	3,527	4,052	4,823
Manufacturers-Foreign	1,900	3,500	2,700	3,100	1,900	1,800
Other Business-US	500	500	514	3,655	5,160	6,012
Other Business-Foreign	0	0	0	30	30	0
Joint Ventures-US	0	0	0	0	500	1,000
Joint Ventures-Foreign	0	0	0	0	500	1,000
Foundations-US	300	100	100	100	100	100
Foundations-Foreign	0	0	0	0	0	0
Total (all categories)	59,989	65,212	67,066	71,283	75,460	84,483
PERCENT, BY FUNDING SOURCE						
Internal	50.25	42.38	53.79	42.55	44.06	44.89
Dept. of Defense	31.33	37.70	35.14	35.76	34.24	32.79
Other Federal Govt	6.50	5.09	7.16	5.61	5.57	4.97
State and Local Govt	0.00	0.00	0.00	0.63	0.00	0.00
Foreign Govt	0.83	0.77	0.89	0.98	0.00	0.00
Manufacturers-US	6.58	7.77	9.46	4.95	5.37	5.71
Manufacturers-Foreign	3.17	5.37	4.03	4.35	2.52	2.13
Other Business-US	0.83	0.77	0.77	5.13	6.84	7.12
Other Business-Foreign	0.00	0.00	0.00	0.04	0.04	0.00
Joint Ventures-US	0.00	0.00	0.00	0.00	0.66	1.18
Joint Ventures-Foreign	0.00	0.00	0.00	0.00	0.66	1.18
Foundations-US	0.50	0.15	0.15	0.14	0.13	0.12
Foundations-Foreign	0.00	0.00	0.00	0.00	0.00	0.00
Total (all categories)	100.00	100.00	100.00	100.00	100.00	100.00

Source: U.S. Dept. of Commerce, BXA/OIRA Sector Survey

TABLE 2. ARTIFICIAL INTELLIGENCE
AS REPORTED RESEARCH PERFORMED BY UNIVERSITIES BY SOURCE OF FUNDING, 1989-1994
(in thousands of dollars)

FUNDING SOURCE	1989	1990	1991	1992	1993	1994
Internal	$131	$135	$156	$116	$155	$160
Dept. of Defense	26,638	28,561	28,506	31,644	34,697	36,625
Other Federal Govt	7,678	8,162	11,183	10,991	10,945	11,346
State and Local Govt	968	1,206	1,091	965	1,262	1,313
Foreign Govt	20	195	130	145	0	0
Manufacturers-US	3,505	3,718	3,925	3,889	3,824	3,621
Manufacturers-Foreign	1,757	1,871	2,085	2,180	2,039	2,069
Other Business-US	1,205	1,212	1,303	1,512	1,650	1,617
Other Business-Foreign	0	0	0	0	0	0
Joint Ventures-US	0	3	0	50	22	43
Joint Ventures-Foreign	0	0	40	40	35	35
Foundations-US	439	4,028	3,873	3,614	3,500	3,500
Foundations-Foreign	0	56	0	590	590	560
Total (all categories)	42,341	49,148	52,291	55,736	58,718	60,889
PERCENT, BY FUNDING SOURCE						
Internal	0.31	0.27	0.30	0.21	0.26	0.26
Dept. of Defense	62.91	58.11	54.51	56.78	59.09	60.15
Other Federal Govt	18.13	16.61	21.39	19.72	18.64	18.63
State and Local Govt	2.29	2.45	2.09	1.73	2.15	2.16
Foreign Govt	0.05	0.40	0.25	0.26	0.00	0.00
Manufacturers-US	8.28	7.56	7.51	6.98	6.51	5.95
Manufacturers-Foreign	4.15	3.81	3.99	3.91	3.47	3.40
Other Business-US	2.85	2.47	2.49	2.71	2.81	2.66
Other Business-Foreign	0.00	0.00	0.00	0.00	0.00	0.00
Joint Ventures-US	0.00	0.01	0.00	0.09	0.04	0.07
Joint Ventures-Foreign	0.00	0.00	0.08	0.07	0.06	0.06
Foundations-US	1.04	8.20	7.41	6.48	5.96	5.75
Foundations-Foreign	0.00	0.11	0.00	1.06	1.00	0.92
Total (all categories)	100.00	100.00	100.00	100.00	100.00	100.00

Source: U.S. Dept. of Commerce, BXA/OIRA Sector Survey

TABLE 3. ARTIFICIAL INTELLIGENCE RESEARCH PROJECTS BY AREA OF RESEARCH (Based on survey reponses by companies and universities for research projects active in 1991-1992)						
RESEARCH AREA	RESEARCH PROJECTS AND PROJECT DOLLAR VALUES					
	AT FIRMS		AT UNIVERSITIES		TOTAL	
	no.	dollar value	no.	dollar value	no.	dollar value
APPLICATIONS	50	15,020,340	45	12,452,000	95	27,472,340
A1-Automated Programming	11	758,167	2	76,667	13	834,833
A2-Expert Data Base Systems	4	319,667	4	493,800	8	813,467
A3-Expert Systems	21	9,143,190	27	9,849,258	48	18,992,448
A4-Game Playing	0	0	0	0	0	0
A5-Not Specified	14	4,799,317	12	2,032,275	26	6,831,592
AUTOMATED REASONING	73	28,252,160	79	14,989,850	152	43,242,010
B1-Distributed Problem Solving	5	1,352,600	6	187,500	11	1,540,100
B2-Fuzzy Logic	12	2,316,573	20	1,458,100	32	3,774,673
B3-General Problem Solving	3	706,667	8	6,229,217	11	6,935,883
B4-Pattern Matching	7	1,494,240	4	845,167	11	2,339,407
B5-Reasoning	18	3,723,900	11	2,142,858	29	5,866,758
B6-Search	6	716,467	10	1,385,733	16	2,102,200
B7-Uncertainty	8	1,539,133	13	1,460,100	21	2,999,233
B8-Not Specified	14	16,402,580	7	1,281,175	21	17,683,755
COGNITIVE MODELING	14	10,207,000	25	7,196,408	39	17,403,408
C1-Decision-Modeling System	4	893,333	8	2,338,783	12	3,232,117
C2-Understanding Systems	1	109,500	13	3,725,450	14	3,834,950
C3-Not Specified	9	9,204,167	4	1,132,175	13	10,336,342
(TABLE CONTINUES NEXT PAGE)						

Note: Research projects often included more than one area of research. In cases where projects involved multiple areas, total dollar amounts were proportioned equally to each area mentioned for any given research project. A total of 408 projects were reported by all surveyed respondents. Sixty percent of the projects (248) involved just one area of research. Another 61 projects involved 2 areas; and 41 more involved 3 areas.)

TABLE 3. (continued) RESEARCH AREA	RESEARCH PROJECTS AND PROJECT DOLLAR VALUES					
	AT FIRMS		AT UNIVERSITIES		TOTAL	
	no.	dollar value	no.	dollar value	no.	dollar value
KNOWLEDGE REP.	48	32,048,487	74	16,011,267	122	48,059,753
D1-Knowledge Representation	15	4,376,817	24	4,561,367	39	8,938,183
D2-Neural Networks	18	6,120,740	33	8,490,850	51	14,611,590
D3-Not Specified	15	21,550,930	17	2,959,050	32	24,509,980
LEARNING	26	33,674,207	49	13,182,608	75	46,856,815
E1-Knowledge Acquisition	4	1,600,373	7	1,637,950	11	3,238,323
E2-Machine Learning	7	6,667,333	17	3,523,550	24	10,190,883
E3-Pattern Recognition	7	7,056,500	17	6,730,683	24	13,787,183
E4-Not Specified	8	18,350,000	8	1,290,425	16	19,640,425
NATURAL LANGUAGE UNDERSTANDING	55	48,858,730	25	11,884,300	80	60,743,030
F1-Machine Translation	2	1,450,000	7	2,054,600	9	3,504,600
F2-Natural Lang. Processing	17	11,586,750	9	1,440;700	26	13,027,450
F3-Speech Recognition	26	18,319,900	9	8,389,000	35	26,708,900
F4-Not Specified	10	17,502,080	0	0	10	17,502,080
PLANNING AND PROBLEM SOLVING	25	6,365,097	21	7,955,975	46	14,321,072
G1-Planning	25	6,365,097	21	7,955,975	46	14,321,072
ROBOTICS AND MACHINE VISION	52	27,708,780	68	47,892,092	120	75,600,872
H1-Robotics	6	1,959,100	26	25,736,667	32	27,695,767
H2-Machine Vision	44	24,501,600	40	16,295,425	84	40,797,025
H3-Not Specified	2	1,248,080	2	5,860,000	4	7,108,080
TOTAL (major areas)	343	202,134,800	386	131,564,500	729	333,699,300

Source: U.S. Dept. of Commerce, BXA/OIRA Sector Survey

APPENDIX D

TABLE 4. PERCENT ARTIFICIAL INTELLIGENCE RESEARCH BY AREA OF RESEARCH (Based on total dollar values for companies and universities shown on Table 11)			
RESEARCH AREA	FIRMS	UNIVERSITIES	TOTAL
APPLICATIONS	7.43%	9.46%	8.23%
A1-Automated Programming	0.38	0.06	0.25
A2-Expert Data Base Systems	0.16	0.38	0.24
A3-Expert Systems	4.52	7.49	5.69
A4-Game Playing	0.00	0.00	0.00
A5-Not Specified	2.37	1.54	2.05
AUTOMATED REASONING	13.98	11.39	12.96
B1-Distributed Problem Solving	0.67	0.14	0.46
B2-Fuzzy Logic	1.15	1.11	1.13
B3-General Problem Solving	0.35	4.73	2.08
B4-Pattern Matching	0.74	0.64	0.70
B5-Reasoning	1.84	1.63	1.76
B6-Search	0.35	1.05	0.63
B7-Uncertainty	0.76	1.11	0.90
B8-Not Specified	8.11	0.97	5.30
COGNITIVE MODELING	5.05	5.47	5.22
C1-Decision-Modeling System	0.44	1.78	0.97
C2-Understanding Systems	0.05	2.83	1.15
C3-Not Specified	4.55	0.86	3.10
KNOWLEDGE REPRESENTATION	15.86	12.17	14.40
D1-Knowledge Representation	2.17	3.47	2.68
D2-Neural Networks	3.03	6.45	4.38
D3-Not Specified	10.66	2.25	7.34
(TABLE CONTINUES NEXT PAGE)			

TABLE 4. (continued)			
RESEARCH AREA	FIRMS	UNIVERSITIES	TOTAL
LEARNING	16.66	10.02	14.04
E1-Knowledge Acquisition	0.79	1.24	0.97
E2-Machine Learning	3.30	2.68	3.05
E3-Pattern Recognition	3.49	5.12	4.13
E4-Not Specified	9.08	0.98	5.89
NATURAL LANGUAGE UNDERSTANDING	24.17	9.03	18.20
F1-Machine Translation	0.72	1.56	1.05
F2-Natural Language Processing	5.73	1.10	3.90
F3-Speech Recognition	9.06	6.38	8.00
F4-Not Specified	8.66	0.00	5.24
PLANNING AND PROBLEM SOLVING	3.15	6.05	4.29
G1-Planning	3.15	6.05	4.29
ROBOTICS AND MACHINE VISION	13.71	36.40	22.66
H1-Robotics	0.97	19.56	8.30
H2-Machine Vision	12.12	12.39	12.23
H3-Not Specified	0.62	4.45	2.13
TOTAL (major areas)	100.00%	100.00%	100.00%

Source: U.S. Dept. of Commerce, BXA/OIRA Sector Survey

TABLE 5. ARTIFICIAL INTELLIGENCE RESEARCH PERFORMED BY COMPANIES BY AREA AND TYPE OF RESEARCH
(Based on survey responses from companies for ongoing research in 1991-1992)

RESEARCH AREA	RESEARCH TYPE (dollar value)			
	BASIC	APPLIED	OTHER	TOTAL
APPLICATIONS	3,898,653	5,317,007	2,344,580	11,560,240
A1-Automated Programming	123,358	381,808	5,000	510,167
A2-Expert Data Base Systems	1,667	313,000	5,000	319,667
A3-Expert Systems	3,148,191	2,015,199	1,202,000	6,365,390
A4-Game Playing	0	0	0	0
A5-Not Specified	625,437	2,607,000	1,132,580	4,365,017
AUTOMATED REASONING	5,688,838	9,208,942	1,791,200	16,688,980
B1-Distributed Problem Solving	35,000	85,000	5,000	125,000
B2-Fuzzy Logic	860,945	479,729	651,200	1,991,873
B3-General Problem Solving	579,167	127,500	0	706,667
B4-Pattern Matching	916,443	577,797	0	1,494,240
B5-Reasoning	1,304,000	903,500	5,000	2,212,500
B6-Search	301,833	58,333	0	360,167
B7-Uncertainty	1,667	76,667	5,000	83,333
B8-Not Specified	1,689,783	6,900,417	1,125,000	9,715,200
COGNITIVE MODELING	2,934,833	6,092,667	880,000	9,907,500
C1-Decision-Modeling System	335,667	362,667	5,000	703,333
C2-Understanding Systems	0	0	0	0
C3-Not Specified	2,599,167	5,730,000	875,000	9,204,167

(TABLE CONTINUES NEXT PAGE)

APPENDIX D

TABLE 5. (continued) RESEARCH AREA	RESEARCH TYPE (dollar value)			
	BASIC	APPLIED	OTHER	TOTAL
KNOWLEDGE REPRESENTATION	7,110,913	13,819,714	1,778,780	22,709,407
D1-Knowledge Representation	1,818,807	1,830,010	0	3,648,817
D2-Neural Networks	1,178,303	2,084,537	646,200	3,909,040
D3-Not Specified	4,113,803	9,905,167	1.132,580	15,151,550
LEARNING	4,374,303	17,121,004	2,172,400	23,667,707
E1-Knowledge Acquisition	884,045	704,829	5,000	1,593,873
E2-Machine Learning	1,566,000	455,133	646,200	2,667,333
E3-Pattern Recognition	130,925	279,375	646,200	1,056,500
E4-Not Specified	1,793,333	15,681,667	875,000	18,350,000
NATURAL LANGUAGE UNDERSTANDING	7,553,850	9,250,650	646,200	17,450,700
F1-Machine Translation	937,500	312,500	0	1,250,000
F2-Natural Language Processing	1,851,500	528,500	0	2,380,000
F3-Speech Recognition	1,348,400	723,400	646,200	2,718,000
F4-Not Specified	3,416,450	7,686,250	0	11,102,700
PLANNING AND PROBLEM SOLVING	811,250	2,087,917	880,000	3,779,167
G1-Planning	811,250	2,087,917	880,000	3,779,167
ROBOTICS AND MACHINE VISION	1,018,750	1,706,250	0	2,725,000
H1-Robotics	375,000	375,000	0	750,000
H2-Machine Vision	375,000	525,000	0	900,000
H3-Not Specified	268,750	806,250	0	1,075,000
TOTAL (major areas)	33,391,390	64,604,150	104,931,60	108,488,70

Source: U.S. Dept. of Commerce, BXA/OIRA Sector Survey

APPENDIX D

TABLE 6. ARTIFICIAL INTELLIGENCE RESEARCH PERFORMED BY UNIVERSITIES BY AREA AND TYPE OF RESEARCH
(Based on survey responses from universities for ongoing research in 1991-1992)

RESEARCH AREA	RESEARCH TYPE (dollar value)			
	BASIC	APPLIED	OTHER	TOTAL
APPLICATIONS	5643705	6154603	0	11798308
A1-Automated Programming	40833	35833	0	76667
A2-Expert Data Base Systems	120000	373800	0	493800
A3-Expert Systems	4239788	5609470	0	9849258
A4-Game Playing	0	0	0	0
A5-Not Specified	1243083	135500	0	1378583
AUTOMATED REASONING	11717532	3241193	0	14958725
B1-Distributed Problem Solving	134500	53000	0	187500
B2-Fuzzy Logic	457527	1000573	0	1458100
B3-General Problem Solving	6190403	38813	0	6229217
B4-Pattern Matching	178500	666667	0	845167
B5-Reasoning	1917242	225617	0	2142858
B6-Search	1089283	296450	0	1385733
B7-Uncertainty	536027	924073	0	1460100
B8-Not Specified	1214050	36000	0	1250050
COGNITIVE MODELING	7121468	43815	0	7165283
C1-Decision-Modeling System	2324418	14365	0	2338783
C2-Understanding Systems	3696000	29450	0	3725450
C3-Not Specified	1101050	0	0	1101050

(TABLE CONTINUES ON NEXT PAGE)

TABLE 6. (continued) RESEARCH AREA	RESEARCH TYPE (dollar value)			
	BASIC	APPLIED	OTHER	TOTAL
KNOWLEDGE REPRESENTATION	10321307	4920268	16000	15257575
D1-Knowledge Representation	4267027	194340	0	4461367
D2-Neural Networks	4026047	4448803	16000	8490850
D3-Not Specified	2028233	277125	0	2305358
LEARNING	6843105	5149937	567000	12560042
E1-Knowledge Acquisition	1637950	0	0	1637950
E2-Machine Learning	2836938	670612	16000	3523550
E3-Pattern Recognition	1806358	4373325	551000	6730683
E4-Not Specified	561858	106000	0	667858
NATURAL LANGUAGE UNDERSTANDING	8171850	3359450	0	11531300
F1-Machine Translation	1701600	0	0	1701600
F2-Natural Language Processing	1411250	29450	0	1440700
F3-Speech Recognition	5059000	3330000	0	8389000
F4-Not Specified	0	0	0	0
PLANNING AND PROBLEM SOLVING	7820110	135865	0	7955975
G1-Planning	7820110	135865	0	7955975
ROBOTICS AND MACHINE VISION	40808598	7047493	0	47856092
H1-Robotics	25670167	66500	0	25736667
H2-Machine Vision	12794432	3500993	0	16295425
H3-Not Specified	2344000	3480000	0	5824000
TOTAL (major categories)	98447675	30052625	583000	129083300

Source: U.S. Dept. of Commerce, BXA/OIRA Sector Survey

TABLE 7. PERCENT BASIC, APPLIED, & OTHER ARTIFICIAL INTELLIGENCE RESEARCH BY AREA (Based on data from Table 13-companies and Table 14-universities)						
RESEARCH AREA	RESEARCH BY FIRMS			RESEARCH BY UNIVERSITIES		
	BASIC	APPLIED	OTHER	BASIC	APPLIED	OTHER
APPLICATIONS	33.7%	46.0%	20.3%	47.8%	52.2%	0.0%
A1-Automated Programming	24.2	74.8	1.0	53.3	46.7	0.0
A2-Expert Data Base Systems	0.5	97.9	1.6	24.3	75.7	0.0
A3-Expert Systems	49.5	31.7	18.9	43.0	57.0	0.0
A4-Game Playing	0.0	0.0	0.0	0.0	0.0	0.0
A5-Not Specified	14.3	59.7	25.9	90.2	9.8	0.0
AUTOMATED REASONING	34.1	55.2	10.7	78.3	21.7	0.0
B1-Distributed Problem Solving	28.0	68.0	4.0	71.7	28.3	0.0
B2-Fuzzy Logic	43.2	24.1	32.7	31.4	68.6	0.0
B3-General Problem Solving	82.0	18.0	0.0	99.4	0.6	0.0
B4-Pattern Matching	61.3	38.7	0.0	21.1	78.9	0.0
B5-Reasoning	58.9	40.8	0.2	89.5	10.5	0.0
B6-Search	83.8	16.2	0.0	78.6	21.4	0.0
B7-Uncertainty	2.0	92.0	6.0	36.7	63.3	0.0
B8-Not Specified	17.4	71.0	11.6	97.1	2.9	0.0
COGNITIVE MODELING	29.6	61.5	8.9	99.4	0.6	0.0
C1-Decision-Modeling System	47.7	51.6	0.7	99.4	0.6	0.0
C2-Understanding Systems	0.0	0.0	0.0	99.2	0.8	0.0
C3-Not Specified	28.2	62.3	9.5	100.0	0.0	0.0
(TABLE CONTINUES NEXT PAGE)						

TABLE 7. (continued) RESEARCH AREA	RESEARCH BY FIRMS			RESEARCH BY UNIVERSITIES		
	BASIC	APPLIED	OTHER	BASIC	APPLIED	OTHER
KNOWLEDGE REPRESENTATION	31.3	60.9	7.8	67.6	32.2	0.1
D1-Knowledge Representation	49.8	50.2	0.0	95.6	4.4	0.0
D2-Neural Networks	30.1	53.3	16.5	47.4	52.4	0.2
D3-Not Specified	27.2	65.4	7.5	88.0	12.0	0.0
LEARNING	18.5	72.3	9.2	54.5	41.0	4.5
E1-Knowledge Acquisition	55.5	44.2	0.3	100.0	0.0	0.0
E2-Machine Learning	58.7	17.1	24.2	80.5	19.0	0.5
E3-Pattern Recognition	12.4	26.4	61.2	26.8	65.0	8.2
E4-Not Specified	9.8	85.5	4.8	84.1	15.9	0.0
NATURAL LANGUAGE UNDERSTANDING	43.3	53.0	3.7	70.9	29.1	0.0
F1-Machine Translation	75.0	25.0	0.0	100.0	0.0	0.0
F2-Natural Language Processing	77.8	22.2	0.0	98.0	2.0	0.0
F3-Speech Recognition	49.6	26.6	23.8	60.3	39.7	0.0
F4-Not Specified	30.8	69.2	0.0	0.0	0.0	0.0
PLANNING AND PROBLEM SOLVING	21.5	55.2	23.3	98.3	1.7	0.0
G1-Planning	21.5	55.2	23.3	98.3	1.7	0.0
ROBOTICS AND MACHINE VISION	37.4	62.6	0.0	85.3	14.7	0.0
H1-Robotics	50.0	50.0	0.0	99.7	0.3	0.0
H2-Machine Vision	41.7	58.3	0.0	78.5	21.5	0.0
H3-Not Specified	25.0	75.0	0.0	40.2	59.8	0.0
TOTAL (all categories)	30.8%	59.5%	9.7%	76.3%	23.3%	0.5%

Source: U.S. Dept. of Commerce, BXA/OIRA Sector Survey

TABLE 8. TOTAL RESEARCH FUNDED BY ALL U.S. GOVERNMENT AGENCIES, U.S. DEPARTMENT OF DEFENSE, NATIONAL AERONAUTICS AND SPACE ADMINISTRATION, AND THE NATIONAL SCIENCE FOUNDATION, 1983-1993
(in millions of dollars)

YEAR	TOTAL U.S. GOVERNMENT	U.S. DEPARTMENT OF DEFENSE	U.S. NATIONAL AERONAUTICS & SPACE ADMIN.	U.S. NATIONAL SCIENCE FOUNDATION
1983	$14,253.53	$3,222.65	$1,544.82	$1,061.99
1984	14,978.77	3,048.53	1,709.20	1,202.82
1985	16,133.42	3,168.28	1,783.60	1,345.58
1986	16,502.20	3,227.26	2,069.00	1,353.34
1987	17,940.52	3,347.56	2,269.25	1,470.51
1988	18,650.15	3,239.12	2,331.58	1,532.77
1989	20,765.53	3,656.12	2,878.81	1,670.40
1990	21,731.13	3,529.23	3,060.69	1,689.54
1991	23,968.38	3,717.91	3,371.16	1,785.22
1992	26,195.45	4,097.99	3,664.08	1,936.46
1993	27,899.31	4,526.83	4,163.23	2,246.93
PERCENT OF TOTAL RESEARCH BY THREE AGENCIES				
YEAR	PERCENT FOR THREE AGENCIES	DEFENSE	NASA	NSF
1983	40.90%	22.61%	10.84%	7.45%
1984	39.79	20.35	11.41	8.03
1985	39.03	19.64	11.06	8.34
1986	40.30	19.56	12.54	8.20
1987	39.50	18.66	12.65	8.20
1988	38.09	17.37	12.50	8.22
1989	39.51	17.61	13.86	8.04
1990	38.10	16.24	14.08	7.77
1991	37.03	15.51	14.07	7.45
1992	37.02	15.64	13.99	7.39
1993	39.20	16.23	14.92	8.05

Source: U.S. National Science Foundation, Survey of Federal Funds for Research and Development

TABLE 9. TOTAL BASIC RESEARCH FUNDED BY ALL U.S. GOVERNMENT AGENCIES, U.S. DEPARTMENT OF DEFENSE, NATIONAL AERONAUTICS AND SPACE ADMINISTRATION, AND THE NATIONAL SCIENCE FOUNDATION, 1983-1993
(in millions of dollars)

YEAR	TOTAL U.S. GOVERNMENT	U.S. DEPARTMENT OF DEFENSE	U.S. NATIONAL AERONAUTICS & SPACE ADMIN.	U.S. NATIONAL SCIENCE FOUNDATION
1983	$6,260.13	$785.62	$617.03	$999.14
1984	7,067.36	847.86	754.50	1,132.34
1985	7,818.68	861.41	750.90	1,261.81
1986	8,153.08	923.92	916.70	1,275.22
1987	8,942.43	907.60	1,013.72	1,371.17
1988	9,473.62	876.92	1,112.71	1,433.19
1989	10,602.01	947.93	1,417.41	1,562.63
1990	11,285.59	947.62	1,636.93	1,586.28
1991	12,170.79	994.15	1,705.58	1,676.21
1992	13,254.06	1,094.44	1,902.73	1,808.62
1993	14,184.33	1,161.83	2,060.15	2,093.57
PERCENT OF TOTAL BASIC RESEARCH FUNDED BY THREE AGENCIES				
YEAR	PERCENT FOR THREE AGENCIES	DEFENSE	NASA	NSF
1983	38.37%	12.55%	9.86%	15.96%
1984	38.69	12.00	10.68	16.02
1985	36.76	11.02	9.60	16.14
1986	38.22	11.33	11.24	15.64
1987	36.82	10.15	11.34	15.33
1988	36.13	9.26	11.75	15.13
1989	37.05	8.94	13.37	14.74
1990	36.96	8.40	14.50	14.06
1991	35.95	8.17	14.01	13.77
1992	36.26	8.26	14.36	13.65
1993	37.47	8.19	14.52	14.76

Source: U.S. National Science Foundation, Survey of Federal Funds for Research and Development

TABLE 10. TOTAL APPLIED RESEARCH FUNDED BY ALL U.S. GOVERNMENT AGENCIES, U.S. DEPARTMENT OF DEFENSE, NATIONAL AERONAUTICS AND SPACE ADMINISTRATION, AND THE NATIONAL SCIENCE FOUNDATION, 1983-1993 (in millions of dollars)				
YEAR	TOTAL U.S. GOVERNMENT	U.S. DEPARTMENT OF DEFENSE	U.S. NATIONAL AERONAUTICS & SPACE ADMIN.	U.S. NATIONAL SCIENCE FOUNDATION
1983	$7,993.39	$2,437.03	$927.79	$62.85
1984	7,911.41	2,200.67	954.70	70.48
1985	8,314.74	2,306.88	1,032.70	83.78
1986	8,349.12	2,303.35	1,152.30	78.11
1987	8,998.10	2,439.96	1,255.53	99.34
1988	9,176.53	2,362.20	1,218.86	99.58
1989	10,163.52	2,708.20	1,461.40	107.77
1990	10,445.53	2,581.61	1,423.77	103.26
1991	11,797.58	2,723.75	1,665.59	109.02
1992	12,941.39	3,003.56	1,761.35	127.85
1993	13,714.97	3,365.00	2,103.08	153.36
PERCENT OF TOTAL APPLIED RESEARCH BY THREE AGENCIES				
YEAR	PERCENT FOR THREE AGENCIES	DEFENSE	NASA	NSF
1983	42.88%	30.49%	11.61%	0.79%
1984	40.77	27.82	12.07	0.89
1985	41.17	27.74	12.42	1.01
1986	42.32	27.59	13.80	0.94
1987	42.17	27.12	13.95	1.10
1988	40.11	25.74	13.28	1.09
1989	42.09	26.65	14.38	1.06
1990	39.33	24.71	13.63	0.99
1991	38.13	23.09	14.12	0.92
1992	37.81	23.21	13.61	0.99
1993	40.99	24.54	15.33	1.12

Source: U.S. National Science Foundation, Survey of Federal Funds for Research and Development

TABLE 11. TOTAL RESEARCH FUNDED (IN CONSTANT DOLLARS) BY ALL U.S. GOVERNMENT AGENCIES, U.S. DEPARTMENT OF DEFENSE, NATIONAL AERONAUTICS AND SPACE ADMINISTRATION, AND THE NATIONAL SCIENCE FOUNDATION, 1983-1993
(in millions of constant 1993 dollars)

YEAR	TOTAL U.S. GOVERNMENT	U.S. DEPARTMENT OF DEFENSE	U.S. NATIONAL AERONAUTICS & SPACE ADMIN.	U.S. NATIONAL SCIENCE FOUNDATION
1983	$20,151.54	$4,556.16	$2,184.06	$1,501.43
1984	20,268.31	4,125.07	2,312.78	1,627.58
1985	21,043.59	4,132.54	2,326.43	1,755.11
1986	20,903.92	4,088.09	2,620.88	1,714.32
1987	22,066.84	4,117.49	2,791.17	1,808.72
1988	22,142.55	3,845.68	2,768.18	1,819.80
1989	23,605.92	4,156.22	3,272.58	1,898.88
1990	23,717.20	3,851.78	3,340.42	1,843.95
1991	25,240.67	3,915.26	3,550.11	1,879.99
1992	26,827.98	4,196.95	3,752.56	1,983.22
1993	27,899.31	4,526.83	4,162.23	2,246.93
TOTAL BASIC RESEARCH BY THREE AGENCIES (in constant $93)				
YEAR	TOTAL U.S. GOVERNMENT	DEFENSE	NASA	NSF
1983	$8,850.53	$1,110.70	$872.36	$1,412.57
1984	9,563.09	1,147.27	1,020.94	1,532.21
1985	10,198.28	1,123.57	979.43	1,645.84
1986	10,327.79	1,170.36	1,161.22	1,615.37
1987	10,999.18	1,116.35	1,246.87	1,686.54
1988	11,247.64	1,041.13	1,321.08	1,701.57
1989	12,052.19	1,077.59	1,611.29	1,776.37
1990	12,317.02	1,034.23	1,786.53	1,731.25
1991	12,816.85	1,046.93	1,796.11	1,765.18
1992	13,574.10	1,120.86	1,948.67	1,852.29
1993	14,184.33	1,161.83	2,060.15	2,093.57

(TABLE CONTINUES NEXT PAGE)

TABLE 11. (continued)				
TOTAL APPLIED RESEARCH BY THREE AGENCIES (in constant $93)				
YEAR	TOTAL U.S. GOVERNMENT	DEFENSE	NASA	NSF
1983	$11,301.01	$3,445.45	$1,311.70	$88.86
1984	10,705.21	2,977.81	1,291.84	95.35
1985	10,845.31	3,008.97	1,347.00	109.27
1986	10,576.13	2,917.73	1,459.66	98.95
1987	11,067.66	3,001.15	1,544.30	122.18
1988	10,894.91	2,804.54	1,447.10	118.23
1989	11,553.72	3,078.63	1,661.30	122.51
1990	11,400.18	2,817.55	1,553.89	112.70
1991	12,423.82	2,868.34	1,754.00	114.80
1992	13,253.88	3,076.08	1,803.88	130.93
1993	13,714.97	3,365.00	2,103.08	153.36

Note: Gross Domestic Product deflator used to translate to constant dollars

Source: U.S. National Science Foundation, Survey of Federal Funds for Research and Development

TABLE 12. COMPUTER SCIENCE RESEARCH AS PERCENT OF TOTAL RESEARCH FUNDED BY ALL U.S. GOVERNMENT AGENCIES, U.S. DEPARTMENT OF DEFENSE, NATIONAL AERONAUTICS AND SPACE ADMINISTRATION, AND THE NATIONAL SCIENCE FOUNDATION, 1983-1993
(in thousands of dollars)

	ALL AGENCIES		DEFENSE		NASA		NSF	
YEAR	AMOUNT	%	AMOUNT	%	AMOUNT	%	AMOUNT	%
1983	214,555	1.51	114,364	3.55	42,349	2.74	34,370	3.24
1984	214,528	1.43	103,114	3.38	52,994	3.10	33,006	2.74
1985	280,270	1.74	143,715	4.54	51,692	2.90	44,432	3.30
1986	302,426	1.83	163,995	5.08	39,745	1.92	45,188	3.34
1987	297,843	1.66	149,720	4.47	45,867	2.02	60,084	4.09
1988	292,714	1.57	154,455	4.77	40,488	1.74	55,301	3.61
1989	364,621	1.76	200,475	5.48	48,148	1.67	71,333	4.27
1990	562,128	2.59	355,210	10.06	45,216	1.48	109,092	6.46
1991	584,941	2.44	288,556	7.76	51,080	1.52	106,226	5.95
1992	778,185	2.97	445,318	10.87	47,069	1.28	118,118	6.10
1993	811,859	2.91	425,618	9.40	57,756	1.39	153,910	6.85
BASIC RESEARCH IN COMPUTER SCIENCE AND PERCENT OF TOTAL BASIC FUNDING								
YEAR	AMOUNT	%	AMOUNT	%	AMOUNT	%	AMOUNT	%
1983	90,441	1.44	45,318	5.77	7,350	1.19	32,377	3.24
1984	104,789	1.48	44,223	5.22	20,856	2.76	32,058	2.83
1985	116,071	1.48	39,674	4.61	18,755	2.50	43,004	3.41
1986	131,401	1.61	39,839	4.31	20,004	2.18	44,099	3.46
1987	128,547	1.44	38,293	4.22	19,052	1.88	57,736	4.21
1988	125,496	1.32	41,897	4.78	17,156	1.54	52,121	3.64
1989	159,544	1.50	54,810	5.78	22,221	1.57	67,301	4.31
1990	225,209	2.00	85,197	8.99	21,902	1.34	96,749	6.10
1991	223,516	1.84	69,775	7.02	24,190	1.42	98,511	5.88
1992	274,553	2.07	96,591	8.83	24,906	1.31	109,825	6.07
1993	312,892	2.21	92,667	7.98	25,533	1.24	143,419	6.85

TABLE CONTINUES ON NEXT PAGE

TABLE 12. (continued)								
APPLIED COMPUTER SCIENCE RESEARCH AS PERCENT OF TOTAL APPLIED FUNDED								
	ALL AGENCIES		DEFENSE		NASA		NSF	
YEAR	AMOUNT	%	AMOUNT	%	AMOUNT	%	AMOUNT	%
1983	124,114	1.55	69,046	2.83	34,999	3.77	1,993	3.17
1984	109,739	1.39	58,891	2.68	32,138	3.37	948	1.35
1985	164,199	1.97	104,041	4.51	32,937	3.19	1,428	1.70
1986	171,025	2.05	124,156	5.39	19,741	1.71	1,089	1.39
1987	169,296	1.88	111,427	4.57	26,815	2.14	2,348	2.36
1988	167,218	1.82	112,558	4.76	23,332	1.91	3,180	3.19
1989	205,077	2.02	145,665	5.38	25,927	1.77	4,032	3.74
1990	336,919	3.23	270,013	10.46	23,314	1.64	12,343	11.95
1991	361,425	3.06	218,781	8.03	26,890	1.61	7,715	7.08
1992	503,632	3.89	348,727	11.61	22,163	1.26	8,293	6.49
1993	498,967	3.64	332,951	9.89	32,223	1.53	10,491	6.84

Source: U.S. National Science Foundation, Survey of Federal Funds for Research and Development

APPENDIX D

TABLE 13. COMPUTER SCIENCE RESEARCH (IN CONSTANT DOLLARS) FUNDED BY ALL U.S. GOVERNMENT AGENCIES, AND THE AMOUNT FUNDED AND PERCENT OF TOTAL BY THE U.S. DEPARTMENT OF DEFENSE, NATIONAL AERONAUTICS AND SPACE ADMINISTRATION, AND THE NATIONAL SCIENCE FOUNDATION, 1983-1993
(in thousands of constant 1993 dollars)

YEAR	TOTAL ALL AGENCIES	MAJOR FUNDING SOURCES			PERCENT BY MAJOR SOURCE			
		DoD	NASA	NSF	TOP 3	DoD	NASA	NSF
1983	303,336	161,687	59,873	48,592	89.06	53.30	19.74	16.02
1984	290,285	139,527	71,708	44,662	88.15	48.07	24.70	15.39
1985	365,570	187,454	67,424	57,955	85.57	51.28	18.44	15.85
1986	383,094	207,738	50,346	57,241	82.31	54.23	13.14	14.94
1987	366,347	184,156	56,416	73,903	85.84	50.27	15.40	20.17
1988	347,527	183,378	48,070	65,657	85.49	52.77	13.83	18.89
1989	414,495	227,897	54,734	81,090	87.75	54.98	13.20	19.56
1990	613,503	387,674	49,348	119,062	90.64	63.19	8.04	19.41
1991	615,991	303,873	53,791	111,865	76.22	49.33	8.73	18.16
1992	796,975	456,071	48,206	120,970	78.45	57.23	6.05	15.18
1993	811,859	425,618	57,756	153,910	78.50	52.43	7.11	18.96
BASIC COMPUTER SCIENCE RESEARCH								
1983	127,865	64,070	10,391	45,774	94.03	50.11	8.13	35.80
1984	141,794	59.840	28,221	43,379	92.70	42.20	19.90	30.59
1985	151,397	51,749	24,463	56,092	87.39	34.18	16.16	37.05
1986	166,450	50.465	25,340	55,862	79.10	30.32	15.22	33.56
1987	158,113	47,100	23,434	71,015	89.52	29.79	14.82	44.91
1988	148,996	49,743	20,369	61,881	88.59	33.39	13.67	41.53
1989	181,367	62,307	25,260	76,507	90.47	34.35	13.93	42.18
1990	245,792	92,983	23,904	105,591	90.52	37.83	9.73	42.96
1991	235,381	73,479	25,474	103,740	86.11	31.22	10.82	44.07
1992	281,183	98,923	25,507	112,477	84.25	35.18	9.07	40.00
1993	312,892	92,667	25,533	143,419	83.61	29.62	8.16	45.84

(TABLE CONTINUES NEXT PAGE)

TABLE 13. (continued)								
APPLIED COMPUTER SCIENCE RESEARCH								
YEAR	TOTAL ALL AGENCIES	MAJOR FUNDING SOURCES			PERCENT BY MAJOR SOURCE			
		DoD	NASA	NSF	TOP 3	DoD	NASA	NSF
1983	175,472	97,617	49,481	2,818	85.44	55.63	28.20	1.61
1984	148,492	79,687	43,487	1,283	83.81	53.66	29.29	0.86
1985	214,173	135,706	42,961	1,863	84.29	63.36	20.06	0.87
1986	216,643	157,273	25,007	1,379	84.77	72.60	11.54	0.64
1987	208,234	137,055	32,982	2,888	83.04	65.82	15.84	1.39
1988	198,531	133,635	27,701	3,775	83.17	67.31	13.95	1.90
1989	233,128	165,590	29,473	4,584	85.64	71.03	12.64	1.97
1990	367,711	294,690	25,445	13,471	90.73	80.14	6.92	3.66
1991	380,610	230,394	28,317	8,125	70.11	60.53	7.44	2.13
1992	515,793	357,148	22,698	8,493	75.29	69.24	4.40	1.65
1993	498,967	332,951	32,223	10,491	75.29	66.73	6.46	2.10

Source: U.S. National Science Foundation, Survey of Federal Funds for Research and Development

CRITICAL TECHNOLOGY ASSESSMENT OF THE
U.S. ARTIFICIAL INTELLIGENCE SECTOR

APPENDIX E

OTHER OIRA ASSESSMENTS

U.S. DEPARTMENT OF COMMERCE
BUREAU OF EXPORT ADMINISTRATION
OFFICE OF INDUSTRIAL RESOURCE ADMINISTRATION
STRATEGIC ANALYSIS DIVISION

DEPARTMENT OF COMMERCE

UNITED STATES OF AMERICA