AN ANALYSIS OF FACTORS DIRECTING THE ADMISSION PROCESS OF ARTIFICIAL INTELLIGENCE TECHNOLOGIES

Vasilis A. Vassalos **Suresh Venkatasubramanian**

Computer Science Department
Stanford University
Stanford, CA 94305
email: {vassalos, suresh}@cs.stanford.edu

ABSTRACT

Our main motive for this paper has been the observation that, while artificial intelligence is a prominent branch of computer science that has even found its way into popular culture, it appears that its "fruits" are an inconsequential part of the computer industry. The paper tries to assess the extent to which this observation is accurate, by determining and analyzing some important factors that affect the fortunes of AI technologies in the marketplace. It identifies the stages in the admission process of Information Technology and the particular obstacles AI is facing during that process. Finally, based upon the patterns that the research revealed, we make some predictions about the commercialization of AI in the future.

INTRODUCTION

> "Digital computers can *readily* be programmed to exhibit modes of behavior [...] usually associated only with the nervous systems of living organisms."

An optimistic prediction of a university professor or research scientist of today? Hardly. The above text is the opening paragraph of a paper written by members of the Computer Laboratory of Harvard University and delivered in the ACM National Meeting in 1952.

The challenging nature of artificial intelligence research is perhaps what makes many people suspicious towards its proposed practical applications: As Professor J.D. Ullman puts it [25]: "The fundamental problem of AI research is that it starts from lofty goals and then retreats. [...] There are two kinds of AI: one deals with really perceived hard problems, like vision and planning [...] The other [areas], fuzzy logic, neural nets, expert systems and the like have very limited usefulness". Others tend to see the problem in slightly different terms: In the words of Harvey P. Newquist [18], author of the *Brain Makers*, a history of the AI industry, "It has become increasingly evident [...] that there are two AIs: one theoretical and one practical." This kind of distinction is one of the main reasons of confusion and misunderstanding about AI applications.

We define AI technology as a tool to extract and synthesize knowledge from information. Under that definition, today's AI systems would

- Help organisations manage knowledge assets and deal with complexity
- Help experts solve difficult analysis problems and design new devices.
- Learn from examples
- Provide answers to English questions using both structured data and free text.

Our focus is on the admission of artificial intelligence technology, as part of a process of transfer of technology. Technology transfer is about delivering "approaches, ideas and skills" [20]. We regard admission of the technology by the organisational user as the most critical step in this process.

In the remainder of this paper, the admission process will be viewed in the context of Information Technology. In the following sections we will specifically concentrate on some of the unique problems that arise in the admission of AI-related technologies by prospective users, discuss current trends and suggest a new direction for artificial intelligence technology.

COMMERCIALIZATION OF AI

The Past

Artificial intelligence started the transition from research area to marketable commodity with the formulation of the Knowledge Principle: "The level of competence in performing an intelligent task is a function of the available knowledge." Knowledge representation and expert systems became the main focus of AI researchers, and within a few years, the technology had reached a stage where the basic theory was understood well enough to try and market it.

A few AI firms had already been founded in the late 70's (Larry Harris founded AICorp as early as 1975). The beginning of the 80's signaled the arrival of commercial AI applications on a larger scale. LMI, Symbolics, IntelliCorp, Teknowledge, Inference Corp and Carnegie Group were all founded between 1980 and

1984, leading to comments about a "California Gold Rush" approach to AI solutions [11]. For the next few years, expert systems and tools began to enter the market - in manufacturing, as a decision support tool, for nuclear test ban verification, in the financial sector, the Department of Defense, and in many other areas.

These systems were efficient, and their use has often increased productivity by factors of 10 or more. DuPont has saved over $1 billion using an expert sys tem that cost $50,000 to develop [15]. The Authorizer's Assistant used by the American Express credit card division was found to be 96.5% accurate. However, the systems were built mostly in LISP, a language no MIS or DP people were familiar with, and required specialized hardware. The installation of an Expert System (ES) required a break with traditional COBOL-based information systems, which was a shift many companies were not inclined to make. In the words of Harry Reinstein, founder of Aion, a successful producer of expert system development tools: "Our biggest competitor was not AICorp, it was COBOL".

Nonetheless, the field grew, and from about 30 fielded systems in 1985, the number went to more than 1100 in 1987. The size of the market grew correspondingly: the AI Industry crossed the $1 billion mark in 1986. But even in those times of growth, AI firms did not have a clear sense of direction: In 1986, Teknowledge was reported to "change its products strategy every quarter". The distribution of revenue among industry segments was also worrisome (fig.1).

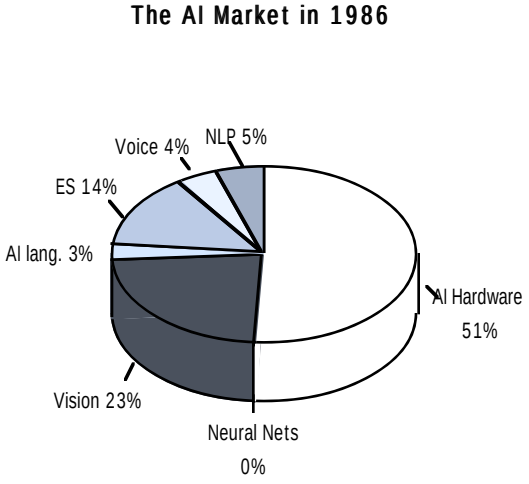

fig.1

As the PC revolution moved into high gear, the AI hardware segment, until then mainly in the business of LISP machines, suffered. The advent of high performance PCs and workstations running LISP environments made stand-alone LISP machines a much more expensive proposition, and companies started folding up.

Specialized AI hardware has since all but disappeared: LMI filed for bankruptcy as soon as 1987 and Symbolics, once the front runner of the entire industry, did the same in 1993, after struggling for many years.

In the corporate world, AI soon became synonymous with broken promises and a lot of hype, In order to be able to sell their "knowledge-based" products, companies attempted to clean the "AI" taint from them, primarily by abandoning LISP in favor of C in the implementation of their products.

The current position

Over 70-80% of Fortune 500 companies are today using some AI technology, and many new systems rely on AI technology, like the AT&T digital PABX, and the spell checker in MS Word. The current installed base of expert systems is now in the order of tens of thousands [5]. A number of new technologies have also appeared in the market in recent years.

One of the newest technologies to come out of the AI labs is the neural net. Neural nets are spreading rapidly in the corporate world, mainly because of their ability to learn, the lack of reliance on the codified knowledge of an expert, and low cost. It is estimated that neural network software sales increased 65% in 1993 to $44 million in 1994 [26]. They are used in pattern recognition systems in banks, for checking credit ratings, and even to predict the stock market.

Fuzzy logic is a another technology that is making its presence felt in the US. The market for fuzzy logic today is mainly in control systems. The GM Saturn uses fuzzy logic in its transmission. Several appliance companies use it in products ranging from refrigerators to microwaves. The size of the market as early as 1991 was reported at $150 million (Cognizer Almanac, 1991).

However, the global AI market (estimated at $900 million in 1993) has been stagnant for the last few years. The overall North American market for AI systems was in 1993 estimated at $601 million [11], with more and more systems employing neural nets and fuzzy logic. Compare this with Oracle Corp., which in 6 years from 1983 to 1989, grew from a $5,000,000 to a $584,000,000 company, and is now estimated at over $1.5 billion. The paradox of AI is that although the technology has apparent and proven utility value, it didn't have that kind of impact on the market.

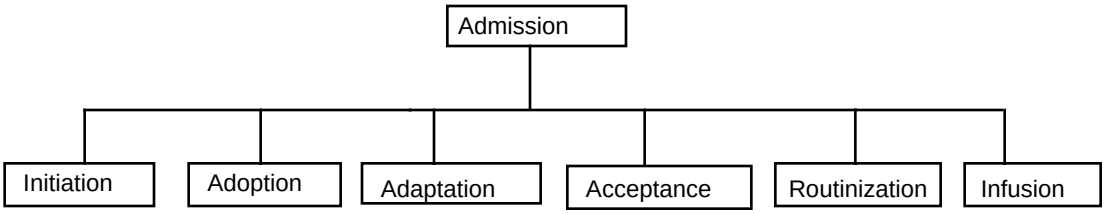

fig.2

ADMISSION PROCESS OF INFORMATION TECHNOLOGY

There have been few proposals for a framework of the admission process of Information Technology (IT). We adopt a general model proposed by Zmud et al. [13],[4]. It is essentially a sequential model of six stages: initiation, adoption, adaptation, acceptance, routinization and infusion. Notice that the process of admission does not stop with the adoption of the technology: the model actually is more focused on what happens *after* that: (fig.2)

- Adaptation

This stage captures the changing nature of technology and acknowledges the incremental changes that accompany its use in a commercial environment. We will not examine that stage.

- Acceptance

Empirical studies confirm that three factors, *attitude towards use, intention to use* and *frequency of use* can capture user attitudes towards a new technology. It seems that there is a link between the user's beliefs about the technology's usefulness and his/her intention to use it. It also seems that managerial intervention generally does not affect the user's attitude toward the technology. That attitude can be changed through increased knowledge of the technology [22]. We will be focusing on attitude toward use and intention to use, in the context of AI technologies.

- Routinization

Routinization refers to the fact that truly successful innovations after a period of time become institutionalized: they become part of the daily work routine of the organisation. It has been established that routinization means three things: *perception of innovation use as "normal", standardized use* and *integration in the administrative infrastructure.* We will look at the success AI technologies had in achieving this status.

- Infusion

Finally, infusion refers to the state where a technology becomes deeper embedded in the organisational structure. The time required for a technology to achieve that status is usually substantial. We will not be looking into this stage, mainly because of the relative youth of commercial AI applications.

The mere adoption of a technology does not guarantee a long-lasting productive presence in the organizational infrastructure. The technology is put to full use when it establishes itself by moving through the admission process.

THE CASE OF ARTIFICIAL INTELLIGENCE

In order to verify our claim that artificial intelligence technologies rarely moved through the admission process, we need to determine the major factors that pertained to AI and adversely affected its acceptance and routinization. In this section we discuss some important issues that explain the fortunes of "intelligent" technologies in the commercial world.

The Image problem

From its early years, Artificial Intelligence has been afflicted by an image problem. The successes that came out of early AI research efforts tended to drift away from the field, and become independent areas of their own, and as a result, AI was perceived as a domain of unsolvable problems. This, however unfortunate, was unavoidable, because of the unclear perceptions we had about the nature of intelligence. As Marvin Minsky points out, "Intelligence is an attribute people ascribe to mental behavior that they admire but do not understand"[23]. That situation sometimes led people to attribute the success of any intelligent-looking application to "mere programming".

Coupled with this was the media hype that accompanied the advent of the AI firms, often fueled by the overoptimistic comments and predictions made by some AI pioneers and visionaries. The overpromise, and a tendency to gloss over the deeper problems (e.g. knowledge acquisition and representation), killed many good applications, and created an environment where, as early as 1984, people were talking about a widespread disdain of AI.

Selling technology, not solutions

One feature of AI that made it different from many other technologies was that the initial impetus for it came not from the industry, but from academia. AI, in other words, was a supply-push technology. As a result of this, there has been a chronic mismatch between what the user expected from AI software, and what the AI firms provided. What the end user needs is a solution to his problem; the nature of that solution was an issue of secondary importance, compared to his estimate of how useful the solution is. Expert systems were sold as technology, while organisations purchased them as an answer to a problem.

It turned out that this answer was sometimes wrong, or even referred to the wrong problem. Medical information systems focused on diagnosis, when the medical community mainly demanded smarter record-keeping. Often, the packages contained only barebones technology that the users were unable and unwilling to tame. This lack of focus on the customer's needs hampered the continued acceptance of the technology. Additionally, there was a lack of expertise in designing expert systems. So, even though the technological aspects of AI may have been revolutionary, entrepreneurial efforts met with early failures.

Marketing Strategies

Dr. Larry Harris, founder of AICorp, has recently said that the success of artificial intelligence companies is going to depend on marketing as opposed to technical differences between the products. The AI industry has not been willing to understand that point.

For a nascent technology, long-term success is crucially dependent on initial performance. Unfortunately, instead of trying to tackle problems which were well understood, and could be easily codified, many of the AI companies touted their products as a panacea for all the troubles of an organisation. The whole industry suffered when the systems could not live up to these promises.

Another marketing mistake was that the AI applications market was initially restricted to R&D departments within large corporations. Focusing on this low-volume, technology-aware segment allowed AI corporations to ignore the "user-unfriendliness" of their software. In the long term, that became a major hurdle in selling their products to the nontechnical end user. What further aggravated the situation was that most of the initial AI entrepreneurs came from an academic setting and so were not corporate "insiders".

Immaturity of the underlying technology

Very little work has been done on “shrink-wrapping” expert system technology — knowledge representation is still more of an art than a science. Moreover, AI was a “fragile” technology: results obtained in a lab were not a good predictor of actual performance in the field. For expert systems, and more so for neural nets, how a specific result is obtained is not clear, mainly because they are declarative and not procedural systems. Therefore, it is difficult to guarantee the performance and the accuracy of the system.

Cultural factors

AI technology, unlike others, tries to simulate *intelligent* behavior, and so it comes into direct conflict with humans performing the same tasks. However, computer systems are not accountable the way humans are. In many systems, AI systems are used to make crucial decisions, decisions concerning loan sanctions, credit card fraud, and stock market investment, among others. As we saw above, these systems are not validated, and so managers feel uneasy about trusting them. AI companies made the mistake of marketing their systems, not as "advisors/helpers", but as "experts", thus alienating many decision-makers that felt uneasy and threatened.

Reluctance to the routinization of expert systems also came from the people who would ultimately be using and maintaining them — the managers in the MIS departments within corporations. Accustomed to mainframes and COBOL, they did not feel comfortable with the LISP/LISP machine environment, and helped influence top management against such systems.

Unfortunate timing

At the time when AI systems first entered the market, widespread computing facilities were still a thing of the future. The PC revolution was just starting to get under way, and it was not yet clear what role computers would play in the corporate world. As a result, selling AI software was much more difficult than it would be now. The global recession and the stock market crash also hindered the growth of AI firms. When the recession struck, AI products and personnel were the first to suffer, because they were the newest and the most suspect costs in an organisation.

Still more problems

Since AI proposes the automation of mental tasks, its benefits are often hard to quantify. It takes a substantial amount of time before such a system can produce a noticeable effect on a company's productivity, and so there is a large gestation period before it can earn the faith of company management.

Last, but not least, is the problem of customer support. Early AI professionals were expensive and in short supply. Their focus therefore was on development of newer and better systems, so they ignored the crucial issue of software maintenance. As a result, customer support for those applications inevitably suffered. It was not uncommon for major, successful expert sys-

tems to slide into obsolescence, mainly due to the lack of regular upgrading and maintenance.

TRENDS AND DIRECTIONS

It is clear from the above analysis that various spin-offs of AI research have failed to establish themselves as industrial strength applications. Nevertheless, there are a couple of positive signs that must be taken into account.

The consolidation of the industry, a direct result of the harsh shakeout that took place in the late 1980s, is bound to prove beneficial to the field. The concentration will also hopefully provide for some long-sought standards.

Another important aspect that implies acceptance of AI technologies is the number and size of in-house projects, that are not included in any industry classification. The bulk of neural nets, fuzzy systems and even knowledge-based systems are developed internally by the government and corporations [5]. It seems that in-house development is a major trend in AI, that explains the fact that many success stories go unannounced. But that trend suggests a new direction for AI technology, one that AI purists might not like.

A new path

There is a new trend growing among AI users, developers and a number of researchers towards acceptance of the fact that AI is often not going to be the definitive solution to the problems that individuals or organisations face; it is generally going to be only a piece in a puzzle.

The key to the acceptance and routinization of AI technology at this point seems to be a low profile. As Prof. Jim Hutchinson (MIT AI Lab) recently said: "Problems of reduced size create high expectations about depth"[1]. Intelligent systems can profit from concentrating in handling more mundane tasks in a more intelligent way. Perhaps the most efficient way to do that is to bring artificial intelligence applications to the domestic user. Several of the problems that arise every day in the home, from plumbing to cooking to gardening could be eased with the use of "simple" AI. The benefits to home users are apparent. The gain to the field would be access to a huge market. Of course, drastic changes in the way such systems are designed will have to be made. Like any consumer software, it must be easily usable. It should also be easy to configure the system to tailor it to a user's specific needs.

That implies a shift in the idea of what intelligent systems do towards what we would probably call "weakest AI": for all practical purposes, expert systems technology should be viewed merely as good software. Once the ghost of "intelligence" has been laid to rest, it will be easier to concentrate on the strengths of expert systems - their ability to provide better leveraging, distribution, and enhancement of valuable expertise.

In today's Information Age, the emphasis is shifting from information processing to knowledge processing. At the same time, we lack an automated tool to help convert information into knowledge. Another problem modern organisations face is that of knowledge distribution. Currently, only a small portion of the organisational knowledge is available at the point of decision, "at the moment of truth". It is here, in the area of knowledge processing and distribution, that AI has a large role to play. Over the past few decades, it has developed many of the tools to process knowledge and provide decision support. Therefore, it is in the best position to corner this vast untapped market. It is estimated that less than 0.1% of current applications can claim to deal with such issues [15].

To take a broader view, we can see AI as part of a growing Knowledge Automation Industry, along with OOP, networks, and CASE [5]. Huge databases are being built, and data warehousing is becoming a central issue in database management. Knowledge based systems will have a large role to play in managing such massive quantities of data.

Another new direction that AI products should follow is in embedded software. We are not alone in speculating that in 10 years, there will be some form of AI in almost all software packages. Similarly, microprocessors using fuzzy logic should become increasingly popular. According to some predictions, by 1996 or 1997 fuzzy microprocessors could amount to half the embedded applications market [27].

Continued Interaction

"The essential technologies to build computers came from major laboratories that were supported by companies that were almost monopolistic. And they brought an enormous constancy of purpose and funding that is never true at a university", James Gibbons, Dean of School of Engineering, Stanford University [9]. For AI, such a base of corporate research support has not been set up yet. And unless industry participates in AI research in a major way, the current mismatch that exists between research and industry needs is going to persist.

That does not mean that industry and academe should drift apart. The problems with AI had their roots exactly in the industry-academe divide, which prevented

[1] *Information and Decision Technologies*, vol 199, 1994

researchers from understanding the nature of the demand for AI applications. On the flip side, MIS professionals were unwilling to use the produced technology in their companies. It is obviously critical to establish connections between the people developing AI systems and the people advancing the state-of-the-art in AI theory; it is still in universities where new ideas will most likely be born.

Short-term growth

Even under the best of circumstances, we do not believe that there will be a significant change in the size of the AI industry in the next few years. We are especially cautious towards a couple of new and promising technologies.

Intelligent agents are the current buzzword in the AI world, a whole issue of the *Communications* of the ACM having been devoted to them [3]. It is highly unlikely that they will be allowed to roam around on the Internet unchecked [25], as they bear an uncanny functional resemblance to viruses. The Internet Worm is one example of what could happen to a rogue agent. Some of the problems encountered in the creation of agents (e.g. queries among different databases) have resisted earlier efforts from researchers in other fields. Finally, the point has been raised that agents will create pockets of over-automation, requiring more tending and feeding than doing a task by hand.

On a similar note, we should put the short-term growth potential of neural nets in perspective. Most industry watchers today rate neural networks as the technology of the future. Neural nets have the distinctive property of being able to adapt to a changing environment, making them favorable in many contexts to a conventional KBS. Nevertheless, it must be added that neural nets are still quite firmly in the experimental stage, and their learning ability is rather limited.

CONCLUSIONS

It has been shown that the limited penetration of artificial intelligent technology in the business community is due to a number of external and internal factors. A sequential model for the admission process of Information technologies is proposed to substantiate the claim that applications of AI have rarely and for good reason moved beyond the adoption stage. It is finally argued that AI systems have gradually made a niche for themselves in the corporate world, and are now well positioned to assume a central role in both corporate and home computing.

ACKNOWLEDGMENTS

The authors would like to thank Avron Barr and Shirley Tessler for helpful comments and numerous fruitful discussions. This work was undertaken as part of the Computer Industry Project at Stanford University.

REFERENCES

[1] Artificial Intelligence Industry Report, DM DATA INC, 1987

[2] Coleman, K. "The AI marketplace in the year 2000", in *AI Expert*, Jan. 1993,. pp. 35-37

[3] Communications of the ACM, March 1994, special issue on AI

[4] Cooper and Zmud, “Information Technology Implementation Research: A Technological Diffusion Approach”, in *Management Science*, vol36, no 2, 1990

[5] Critical Technology Assessment of the U.S. Artificial Intelligence Sector, U.S. Department of Commerce, Office of Industrial Resource Administration, August 1994

[6] Feigenbaum, Edward A., The Rise of the Expert Company, Times Books, 1989

[7] Feigenbaum, E.A., Tiger in a Cage, Video Presentation, 1993

[8] Feigenbaum, E.A., Private communication, 11/94.

[9] Geppert, Linda, "Industrial R&D: the new priorities", in *IEEE Spectrum*, September 1994

[10] Hecht, Lee, "The AI Triangle: R&D, End user and Mainstream markets", in *AI Expert*, December 1986

[11] *Intelligent Software Strategies*, No 1, Jan 1994

[12] King, J.A, "AI gardening", in *AI Expert*, Aug 1994

[13] Kwon and Zmud, “Unifying the fragmented models of Information Systems Implementation”, in *Critical Issues in Information Systems Research*, 1987

[14] Liebowitz, J. "Artificial intelligence and the corporate environment", in *Industry executive* (1989), vol.2, no. 1, pp. 56-58.

[15] Mahler, Ed., Private communication, 11/94.

[16] McEnaey, M, “Teknowledge Retools Expert Systems for Business Market”, in *Computer-World*, Aug 4, 1986

[17] Morris, Charles and Ferguson, Charles, "How architecture wins technology wars", in *Harvard Business Review*, March-April 1993

[18] Newquist, Harvey P. III, "A Tale of two AIs", in *AI Expert*, November 1993

[19] Newquist, Harvey P. III, "AI Burnout", in *Computerworld*, April 18, 1994

[20] Owen, Kenneth, "Technology Transfer as a business", in *Computer Bulletin*, Oct 1990, vol 2, pp. 21-3

[21] Reinstein, Harry, Private communication, 11/94

[22] Saga and Zmud, “The nature and determinants of IT acceptance, routinization and diffusion”, in *Diffusion, Transfer and Implementation of Information*, ed Levine, 1993

[23] Sheil, Beau, "Thinking about Artificial Intelligence", in *Harvard Business Review*, July-August 1987

[24] Tabandeh, Amir "An analysis of artificial intelligence technology in the context of international technology transfer", in *Proceedings, 6th Inter-national Symposium on Artificial Intelligence*, vol. 50, no. 6, pp. 217-222

[25] Ullman, J. D., Private communication, 10/94

[26] US Industrial Outlook, “Computer Software and Networking”, 1994

[27] Walsh, Birrell, "Fuzzy Logic: the American Market", in *AI Expert*, Nov 1993