

If It Works, It's Not AI:

A Commercial Look at Artificial Intelligence Startups

by

Eve M. Phillips

Submitted to the Department of Electrical Engineering and Computer Science

in Partial Fulfillment of the Requirements for the Degrees of

Bachelor of Science in Computer Science and Engineering

and Master of Engineering in Electrical Engineering and Computer Science

at the Massachusetts Institute of Technology

May 7, 1999

[May 7, 1999]

© Copyright 1999 Eve M. Phillips. All rights reserved.

The author hereby grants to M.I.T. permission to reproduce and distribute publicly paper and electronic copies of this thesis and to grant others the right to do so.

Author _____

Department of Electrical Engineering and Computer Science

May 7, 1999


Certified by _____

Patrick Winston
Thesis Supervisor

Accepted by _____

Arthur C. Smith
Chairman, Department Committee on Graduate Theses

ENG



If It Works, It's Not AI:
A Commercial Look at Artificial Intelligence Startups
by
Eve M. Phillips

Submitted to the
Department of Electrical Engineering and Computer Science

May 7, 1999

In Partial Fulfillment of the Requirements for the Degree of
Bachelor of Science in Computer Science and Engineering
and Master of Engineering in Electrical Engineering and Computer Science

ABSTRACT

The goal of this thesis is to learn from the successes and failure of a select group of artificial intelligence (AI) firms in bringing their products to market and creating lasting businesses. I have chosen to focus on AI firms from the 1980s in both the hardware and software industries because the flurry of activity during this time makes it particularly interesting.

The firms I am spotlighting include the LISP machine makers Symbolics and Lisp Machines Inc.; AI languages firms, such as Gold Hill; and expert systems software companies including Teknowledge, IntelliCorp, and Applied Expert Systems. By looking at the history of the technology and international activity around it, such as the Japanese Fifth Generation Project, and profiling a number of firms in these areas, more than enough material is present to offer conclusions that could be relevant to any new technology industry.

From their successes, we can see examples of positive methods of introducing new technologies into the marketplace. However, my choice of this time-period and industry was due to the high level of hype, from the mainstream press and corporate excitement, prior to its downfall. The negative examples of technology business methods from the AI industry offer many more useful lessons. The pitfalls that many of these firms fell into include management inexperience and academic bias, business models which confused products and consulting, misunderstanding their target market, and failing to manage customer and press expectations. These problems are seen as much today as during the lifetime of these companies in high tech markets. While today's high tech firms seem generally better able to understand their market, they still often make similar mistakes. By looking at many situations in which firms faltered, I hope to provide some warnings and suggestions for any company trying to build a business around a new technology.

Thesis Supervisor: Patrick Winston
Title: Professor of Electrical Engineering and Computer Science

Acknowledgements

I would like to thank Patrick Winston and Ed Roberts for their advice, support and encouragement in creating this document. Both provided a level of expertise that was invaluable; without them, I would not have undertaken this project. My thanks are also due to both MIT's Artificial Intelligence Laboratory and the Sloan School of Management, whose members, especially Randy Davis, Tom Knight, Marvin Minsky, Scott Shane, Simon Johnson, Ken Morse and Michael Cusumano, freely supplied many of the stories and details that became the content of this thesis.

My thanks go to all the people outside of MIT who offered their time to tell me about their various AI exploits. For suggesting materials to get me started, my thanks go to Carol Hamilton at the AAI, Mike Brady, and Joe Hadzima. Special thanks to Ken Ross, Douglas Keith, Bob Fondiller, John Sviokla, Chuck Williams, John Tannenbring, Fred Luconi, John Price, Russ Siegelman, Phil Cooper, Ramanathan Guha, Bill Mark, Qunio Takashima, and Judith Bolger for their comments and insight about their experience with the industry. For the unique perspective of the industry press, I would like to thank Harvey Newquist and Don Barker for their help. And to balance out my MIT background, my gratitude goes to Hubert Dreyfus at U.C. Berkeley and Ed Feigenbaum at Stanford University for their interest and ideas.

Finally, I would like to thank my family and friends for their support and help, in particular Joost Bensen and Pam Elder, for their suggestion of contacts, materials, and frameworks for thinking about this thesis; without them this document would have been much less interesting.

Table of Contents

1	INTRODUCTION	6
1.1	BACKGROUND OF AI TECHNOLOGY AND INDUSTRY	6
1.2	SUMMARY OF THESIS POINTS	7
1.3	RESEARCH METHODOLOGY	8
1.4	STRUCTURE OF THESIS CONTENT	8
2	RELATED WORK	10
2.1	THE INTERNATIONAL VIEW: JAPAN AND THE UNITED KINGDOM	10
2.2	COMMERCIALIZING AI	11
2.3	TECHNOLOGY BUSINESS MODELS	13
2.4	FOCUS OF RESEARCH	15
3	WHY THE AI INDUSTRY STUMBLER	16
3.1	MANAGEMENT INEXPERIENCE AND ACADEMIC BIAS	16
3.2	BUSINESS MODELS: PRODUCTS VERSUS CONSULTING	18
3.3	MISUNDERSTANDING THE TARGET MARKET	18
3.3.1	<i>Incompatibility with Clients' Internal Systems</i>	19
3.3.2	<i>Selling Tools Versus Vertical Solutions</i>	20
3.3.3	<i>Selling Technology Versus Real Products</i>	21
3.3.4	<i>Hardware Insufficiency and Cost; Moore's Law</i>	21
3.3.5	<i>Missing the PC Trend</i>	22
3.4	FAILING TO MANAGE EXPECTATIONS	22
4	THE AI SOFTWARE INDUSTRY	25
4.1	BACKGROUND OF EXPERT SYSTEMS	25
4.1.1	<i>Early Expert Systems</i>	27
4.1.2	<i>Issues in Building Expert Systems</i>	27
4.2	THE GANG OF FOUR	28
4.2.1	<i>Carnegie Group Inc.</i>	29
4.2.2	<i>IntelliCorp Inc.</i>	31
4.2.3	<i>Inference Inc.</i>	34
4.2.4	<i>Teknowledge Inc.</i>	37
4.3	EXPERT SYSTEMS SHELLS	39
4.3.1	<i>Artificial Intelligence Corporation</i>	39
4.3.2	<i>Neuron Data, Inc.</i>	41
4.4	SOFTWARE TOOLS	42
4.4.1	<i>Gold Hill Inc.</i>	43
4.4.2	<i>Lucid Inc.</i>	44
4.5	EXPERT SYSTEMS APPLICATIONS	46
4.5.1	<i>Applied Expert Systems Inc.</i>	46
4.5.2	<i>Palladian Software Inc.</i>	49
4.6	KNOWLEDGE BASE FIRMS	51
4.6.1	<i>Cycorp</i>	51
4.7	AI SOFTWARE INDUSTRY SUMMARY	52
5	THE AI HARDWARE INDUSTRY	57
5.1	LISP MACHINES INC.	58
5.2	SYMBOLICS, INC.	60
5.3	OTHER AI HARDWARE PLAYERS	64
5.4	AI HARDWARE INDUSTRY SUMMARY	65
6	CONCLUSIONS AND RECOMMENDATIONS	68

6.1	WHERE ARE THEY NOW?.....	68
6.2	MODERN AI FIRMS.....	69
6.2.1	<i>Silicon Spice</i>	69
6.2.2	<i>i2 Technologies</i>	70
6.2.3	<i>Trilogy Software</i>	71
6.2.4	<i>Ascent Technology</i>	72
6.3	LESSONS FOR HIGH-TECH ENTREPRENEURSHIP.....	74
APPENDIX I: BIBLIOGRAPHY		76
APPENDIX II: LIST OF INTERVIEWS.....		81

1 Introduction

Artificial intelligence (AI) covers many research areas, including robotics, vision systems, natural language, and expert systems. My research focuses mainly on the commercialization attempts in expert systems and the associated hardware (namely Lisp machines, which were optimized for running the Lisp programming language), as well as some natural language software. This focus is due to the large number of companies started in those particular fields during the 1980s and the lessons that can be gleaned from those firms.

The goal of this thesis is to learn from the successes and failures of these technology firms in bringing their products to market and creating lasting businesses. From their successes, we can see examples of positive methods of introducing new technologies into the marketplace. However, my choice of this time-period and industry was due to the many failures as compared to successes, to the point where participants commented that if a program was successful, they would no longer even call it artificial intelligence. The negative examples that these situations provide offer many useful lessons as well. The pitfalls that many of these firms fell into are as true today as during the lifetime of these companies, and while today's high tech firms seem generally better able to understand their market, the same problems still occur. By looking at many situations in which firms faltered, I hope to provide some warnings for any company trying to create a business from a new technology.

1.1 Background of AI Technology and Industry

Because much of the technology for AI firms came from university laboratories (especially the Massachusetts Institute of Technology, Stanford University, and Carnegie Mellon University), it is necessary to look at the history and output of the laboratories to understand the roots of AI commercialization. Serious AI study began in 1956 with the Dartmouth Conference, organized by John McCarthy, then at Dartmouth, Marvin Minsky (then at Harvard), Claude Shannon of Bell Labs, and Nathaniel Rochester of IBM. This conference brought together those who were starting to do work in AI, though little actually resulted from the conference other than coining the term "artificial intelligence". In 1959, Minsky and McCarthy founded the MIT Artificial Intelligence Laboratory, and then in 1962,

McCarthy moved on to found Stanford's AI Laboratory. Thus were two of the centers created which would, in the following decades, develop much of the technology for future AI businesses.¹

Through the 1960s the MIT AI Laboratory received considerable ARPA (Advanced Projects Research Agency, later renamed *Defense* Advanced Projects Research Agency, or DARPA) funding from the government. Stanford and CMU began to build up their own AI laboratories. By the 1970s AI scientists were writing serious AI programs. A lot of money flowed into voice recognition in particular. However, by the late 1970s DARPA funding for AI research began to decline, as AI research had not produced technologies that were clearly useful for military applications. DARPA also changed its funding strategies in the 1970s from funding institutions or individuals to funding specific projects, which also made it more difficult for AI projects to receive funding.² American AI researchers began to look to industry in order to continue their work.³

Other countries also used government and university resources to create industrial AI products. In 1979 Japan organized a meeting to discuss its plans for high technology for the next decade. This meeting paved the way for the joint industry, university, and government supported Fifth Generation Project, starting in 1981. This project focused Japan's efforts on making AI a reality by 1992. Britain responded with the Alvey Project, a project similarly focused on strategic computing. Back in the U.S., AI researchers began to start research centers in large corporations and form independent companies to take advantage of what they saw as great promise in AI. The fate of these companies is the topic of this thesis.⁴

1.2 Summary of Thesis Points

I have attempted to uncover the reasons why the AI "bubble" of the 1980s occurred as it did. The main reasons I have found for this saga, which the rest of this thesis will explain in more detail, are:

- Management inexperience and academic bias from the founders
- Faulty business models: confusion over products versus consulting
- Misunderstanding the target market, including:
 - Incompatibility with clients' internal systems
 - Selling tools (that customers did not have the expertise to use) versus vertical solutions

- Selling technology versus real products to “cross the chasm”
- Hardware insufficiency and cost sensitivity; Moore's Law (as adapted to PCs) and its ramifications for the specialized chip industry
- Missing the PC trend in the corporate market
- Failing to manage expectations of the press and customers

The rest of this thesis will explore these ideas further and give examples of the companies that illustrated these concepts unfortunately much too well.

1.3 Research Methodology

I have relied on a multitude of first and second-hand sources for the information for this thesis. Much of the data came from news articles and press releases chronicling the industry, which I have uncovered through various online news archives, especially Dow Jones Interactive. I have also completed over twenty interviews with individuals involved with various firms (see Interviews section at end) as business people, technologists, advisors, or consultants; my thanks goes to each of them for their time and insights. These interviews were conducted either in person, over the phone, or via email. I have also made use of several books about the industry, most notably Harvey Newquist's *The Brain Makers*, for its data on the artificial intelligence saga.

1.4 Structure of Thesis Content

Section 2 of this thesis reviews related work on the topic of high-tech entrepreneurship and the industry of artificial intelligence. In Section 3, I discuss my major thesis points on the reasons the AI industry stumbled. Section 4 covers the artificial intelligence software industry in particular, and chronicles several different firms in each of the chosen AI subfields, focusing mainly on expert system shells, AI languages, and expert system applications. Section 5 reviews the AI hardware industry, and looks at the chipmakers Lisp Machines Inc. and Symbolics in particular. Section 6 reviews my conclusions and looks at a few of today's promising AI companies and compares their strategies to those of the previous decade.

¹ HP Newquist, *The Brain Makers: Genius, Ego and Greed in the Quest for Machines that Think* (Indianapolis, Sams Publishing, 1994), 45-75.

² Interview with Patrick Winston, 11/23/98.

³ Newquist, *The Brain Makers*, 76-151.

⁴ Newquist, *The Brain Makers*, 151-153.

2 Related Work

Much work has been done on the study of Artificial Intelligence, and also on the topic of high-tech entrepreneurship. In the intersection of the two topics, where this thesis lies, several notable works have been published, although the time at which they were written bears heavily upon the relevance and focus of the work.

From these works I have drawn several preliminary conclusions. Firstly, brilliant AI researchers did not necessarily make successful entrepreneurs. Secondly, the AI community often failed to temper the expectations of their marketplace with the actual current capabilities of AI technologies. Thirdly, despite its stumbles, the entire developed world was clearly fascinated, and still is, by the potential of AI.

2.1 The International View: Japan and the United Kingdom

Many of the works on Japan's Fifth Generation Project outline the possible commercial uses of AI. The original work in this field, *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, by Edward Feigenbaum and Pamela McCorduck, was originally published in 1983 and opened the eyes of the rest of the world to Japan's impending challenge to the AI brain trust. Feigenbaum, then a professor of Computer Science at Stanford University, and also a co-founder of 1980s AI firms Teknowledge and IntelliCorp, aimed to scare his readers into action with inflationary prose about Japan's AI plans. He succeeded in raising interest and directing funds into AI research, though the fears about Japan's imminent superiority in the field turned out to be unfounded. Many American AI firms, hearing about Japan's movements in the field, saw the Project as a serious competitive threat.

Several books written after the project started, including J. Marshall Unger's *The Fifth Generation Fallacy: Why Japan Is Betting Its Future on Artificial Intelligence*, published in 1987, and Michael Cusumano's *Japan's Software Factories*, published in 1991, discuss the (hypothesized) real reasons the Japanese started the project, and why it did not accomplish all of its goals. Unger's book suggests that the Japanese writing system was the main reason behind its support of the Project: Western-designed machines did not handle Japanese characters well. However, as Japan's Project did not succeed as planned, these

authors theorize that the causes of failure included overly high expectations, cultural challenges to setting up this new type of research facility, and lack of enthusiasm among the research community.¹

Looking at the other side of the world, in Britain, several works chronicle the British AI academic and industrial history. The first work to be mentioned must be Sir James Lighthill's report, submitted to the British government in 1973. This report strongly criticized the work being done at the various British university AI labs, and resulted in the Science Research Council cutting AI funding. This took Britain into its own "dark ages" of AI, from which it took ten years to recover, either in a commercial or university setting.

When AI did come back to Britain, it did so with strength, as discussed in the book by Brian Oakley and Kenneth Owen, *Alvey: Britain's Strategic Computing Initiative*. The Alvey Program, started in 1983, was Britain's response to Japan's Fifth Generation Project. This program combined government and industry funding to both educate the market and develop intelligent systems.² The research areas the program focused on were Knowledge-Based Systems, VLSI, Integrated Circuits, Software Engineering, and Speech Technology. However, as the business and cultural climate in Britain is not as amenable to starting companies as the U.S., the technologies were not as commonly commercialized in startup firms.

2.2 Commercializing AI

Harvey Newquist's thorough book on the history of AI and its business applications, *The Brain Makers: Genius, Ego and Greed in the Quest for Machines that Think*, appeared in 1994. This book closely examines the personalities behind much of the AI phenomenon, including people both from the research labs and the companies. This book chronicles the story of many of the early AI firms, but does not analyze too deeply the reasons it happened as it did. This thesis will attempt a more in-depth analysis as well as use a more technical approach to understanding the companies in question.

The historical book *Computer: A History of the Information Machine*, by Martin Campbell-Kelly and William Aspray, covers the more general history of computing, but still provided useful material on the surrounding computing industry that the AI field was participating in.

John Sviokla's 1986 doctoral thesis in Business Administration from Harvard, *PlanPower, XCON, and MUDMAN: An In-Depth Analysis into Three Commercial Expert Systems in Use*, proposes that while an expert system can provide strategic advantages to the firm that uses it, they are still "high-risk, high-

technology ventures which create management problems.”³ His thesis examines the organizational effects of PlanPower (from Applied Expert Systems), XCON (from DEC), and MUDMAN (from CMU and N.L. Baroid) on the companies that use them. His conclusions are that AI, as applied in expert systems, has commercial viability, that the hardware and software tools are powerful enough to do interesting things, and that expert systems can provide firms with a competitive advantage.

In 1983 MIT’s Industrial Liaison Program sponsored a colloquium on applications of AI in business, and brought together speakers with AI interests from academia, finance, and industry, as well as end users of AI. In *The AI Business: The Commercial Uses of Artificial Intelligence*, Patrick Winston and Karen Prendergast (ed.) compiled the speeches from the conference given by all four of the viewpoints represented. From the business aspect, two participants seemed especially applicable. An essay by William Janeway, an investment banker, states that “Only some pieces of the future of Artificial Intelligence should be financed [...], and those may be the ones that by definition no longer are Artificial Intelligence.”⁴ Venture capitalist Frederick Adler described AI as “faddish”, and questioned what the needs were that AI would fulfill. All the participants were generally optimistic about the future of AI, but the immediate financial prospects for the firms creating products using it were unclear.

Philip Cooper’s 1984 Master of Science thesis at the MIT Sloan School of Management, entitled *Artificial Intelligence: A Heuristic Search for Commercial and Management Science Applications*, focuses more on AI applications. Cooper used this thesis as the basis of founding Palladian, a software company that produced corporate advisory software using AI technology. The thesis covers the intellectual history of AI, technical areas of research, and possibilities for commercialization. Cooper recommends using AI in situations in which there is a given domain of knowledge and clear methods to solve the problem.

Startup, by Jerry Kaplan, tells the story of GO Corporation. Kaplan founded GO in 1987 with the plan to build a pen computer using several AI technologies, primarily handwriting recognition. The company did not succeed, and Kaplan’s book gives a detailed analysis of the firm’s rise and fall. In the end, the venture capitalists and the founders were so enthralled by their product that they failed to pay attention to their market’s needs, much like many of the firms described in this thesis.

Unlike most of the other technologists who were writing about AI, Herbert and Stuart Dreyfus criticized “the pretensions of AI and expert systems.” The Dreyfus brothers wrote their book, *Mind Over*

Machine: The Power of Human Intuition and Expertise in the Era of the Computer, in 1986. They state that human intelligence cannot be replicated in a machine, since a machine's way of "thinking" is too fundamentally different from that of a person; in particular, computer processing is far too structured. They claim that computers operate at merely the "rule-following stage," whereas humans surpass this stage and are capable of higher levels of thinking. In respect to AI businesses, they suggest that while there are probably some rule-based functions that computers could replicate, true intelligence can never be copied, and thus much of the hype that swept up the AI industry was unnecessary.

In the second half of the 1980s, many books were published about applying artificial intelligence to business problems. Among these is *Putting Artificial Intelligence to Work: Evaluating and Implementing Business Applications*, by Sy Schoen and Wendell Sykes, published in 1987. One of the authors was from Arthur D. Little's AI Center, and the other was from Litton Industries, where he was an AI program manager. Without mentioning any particular firms, they review the types of problems that they considered AI to be best at solving, and discuss how to manage the process of building an AI solution, whether done in-house or through outside firms. In general, the book promotes a positive view of using AI techniques to solve various problems.

2.3 Technology Business Models

On a more general note, Gordon Bell's *High-Tech Ventures*, published in 1991, makes several mentions of AI technologies in his book for those involved in the high tech world. His approach "enables users to examine all the critical dimensions that affect a new venture."⁵ He claims AI suffered from having a technology but not a product, thus not satisfying any real need. He also critiques some AI firms for trying to establish technical monopolies and for not realizing the time, patience, and capital required to build their market.

Another particularly interesting model comes from Geoffrey Moore's *Crossing The Chasm*. In this book, Moore discusses what separates the successful technology companies from the rest: the ability to "cross the chasm" from the company's early market, dominated by "visionary" customers, into the larger mainstream "pragmatic" market. In fact, Moore singles out AI as one of his examples as a technology that garnered a lot of press and support from its early customers, but never made it into the mainstream. AI suffered from "lack of mainstream hardware, inability to integrate it easily into existing systems, no

established design methodology, and a lack of people trained in how to implement it.”⁶ AI thus fell into two of Moore’s chasms that separate the early and mainstream markets: they took a greater interest in their technology than the industry, and they failed to recognize the importance of the existing infrastructure.⁷ But in terms of Moore’s strategies for crossing the chasm, AI may have had it right; Moore advocates marketing the technology as a radical productivity improvement on some critical success factor of the customer, and that was how many expert systems in particular positioned themselves.⁸ However, this proposition clearly was not enough to make up for AI’s many other problems.

In his book *Entrepreneurs in High Technology*, Edward Roberts describes a series of characteristics that affect a company’s success. First there is the background of the entrepreneur: family background, education, age and work experience (technical, sales, managerial), and personality and motivation. While no single profile fits all entrepreneurs, there were some statistics that indicated that certain profiles are more likely to be successful. A high tech entrepreneur is likely to have a father who was self-employed, a masters degree in engineering, and have an "inventor" personality with a low need for affiliation and a heavy orientation towards independence. Many of the AI entrepreneurs fit this general description, but were heavier on research than development in their previous experience.⁹

At the founding of the company, Roberts found two general factors for success: a strong technological base, in the degree of technology transfer from the source organization, as well as product orientation; and a strong financial base in initial capital. While the AI companies generally had a very high degree of technology transfer, several, especially on the software side, lacked in their product orientation. The range of financial backing varied for most of the firms, but a lack of capital did not seem to be a major problem for most of them.¹⁰

In the next stage, which Roberts calls Postfounding, the company needs to focus on its marketing orientation, namely market interactions and marketing organization and practices; subsequent financing; and managerial orientation, in particular managerial skills acquisition and problem focus. While most of the AI firms built in marketing organizations, they were not very successful in understanding the needs and requirements of their customers. The AI firms also suffered from inexperienced management and an overly academic background. Clearly Roberts' research correlates strongly with the evidence from the early AI industry.¹¹

Indeed, many of these issues were problems for the early AI industry. In this thesis I examine these and other reasons more thoroughly to explain why and how the industry acted as it did.

2.4 Focus of Research

Clearly, from the section above, much varied work has been done by those studying the business of AI and the fate of the firms that attempted it. However, there are several aspects to this research that my thesis will approach differently. To begin, most of the work has focused on particular products, not companies. Some was written before there was enough time after the fact to clearly analyze the subject; in others, the authors lacked a technical background to fully explain the technical issues at stake for the firms involved. Finally, several texts were not focused on the AI industry as a whole, but instead looked just at particular areas (such as expert systems) or firms.

In this thesis, focusing primarily but not exclusively on the decade of the 1980s and the AI companies active at that time, I look at both the technical and business issues that these companies faced. From there, I determine which of those issues were more responsible for the success or failure of the individual firms, as well as for the collapse of the general media opinion of the industry as a whole. I also look at selected firms from before and after the 1980s timeframe in order to make comparisons and to capitalize on the hindsight that writing this thesis now allows.

¹ Michael Cusumano, *Japan's Software Factories*.

² Brian Oakley and Kenneth Owen, *Alvey: Britain's Strategic Computing Initiative*.

³ John Sviokla, *PlanPower, XCON, and MUDMAN*, vii.

⁴ Patrick Winston, *The AI Business*, 271.

⁵ Gordon Bell, *High-Tech Ventures*, v.

⁶ Geoffrey Moore, *Crossing the Chasm* (HarperBusiness 1991), 22-23.

⁷ Moore 57-59.

⁸ Moore 102-104.

⁹ Edward Roberts, *Entrepreneurs In High Technology: Lessons from MIT and Beyond* (Oxford University Press 1991), 245-308.

¹⁰ Roberts.


¹¹ Roberts.

3 Why the AI Industry Stumbled

AI companies faced the same types of issues as most technology companies, as well as most startup firms in any industry. Unfortunately for the participants, the AI industry managed to illustrate unusually well many of the lessons of what not to do when trying to build a business around a new technology. I have described below in some detail the major problems that this thesis is examining. Some of these issues are universal, such as problems in moving from academia to the corporate world; others are more specific to the time and nature of the AI industry, such as the trend towards PCs. All, however, can nevertheless be instructive in understanding technology industries.

How large was the AI industry in the 1980s? Accurate numbers are extremely difficult to come by; very reputable magazines will have very different numbers for the same year. Table 3.1 gives a very general approximation for the industry size; these numbers average a variety of publications and should only be taken as a rough gauge.

Table 3.1: Approximate AI Industry Revenues¹



3.1 Management Inexperience and Academic Bias

Many of these artificial intelligence companies were founded by the creators of the technology, who were often working in academia. While a few academics, like Amar Bose of Bose Corporation, have gone on to found successful companies, that switch is often difficult, as the requirements for success in academia differ greatly from those in the corporate world. The goal of professors and students in electrical

engineering and computer science departments is to produce top-notch technology. A simplified version of how this goal is achieved is that groups within the department obtain contracts, often from DARPA or other government agencies, and occasionally (and increasingly frequently) from corporate sponsors. They then research and produce working demonstrations of hardware, code or combinations of both for these sponsors. The result can reach the marketplace in many different ways: sponsors may take the result and commercialize it; people within the lab may take it to market; or the lab may license it to an existing firm.

While this system works reasonably well for academia, corporate research and development functions quite differently. First, the end product must be a production-ready piece of hardware or software, not just a prototype. A different type of engineering goes into building end-user products, including quality control and manufacturability, that are not as strong considerations in developing the prototype. Academics entering companies must learn this next step to survive. Secondly, money for research and development must come from either funding sources (venture capital, corporate equity agreements, etc.) or sales of existing products; both of which require different methods than winning government contracts; for example, venture capitalists and DARPA have very different goals and must be persuaded with different tactics.

The biggest difference, however, is that while in academia, the goal is to improve the current state of knowledge by creating new technologies using new ideas, in business the goal is profit. While profit can be gained by selling new technologies into a market that can use that technology to solve their problems, the technology itself is only a small, and not even always necessary, part of what makes a successful business. Selling a product is more important than making the product use some new technology or be on the "cutting edge". The technology itself must also be encased in a product that solves a particular problem of the customer, and the marketing of that product must reflect the problem solution, not the technology. Coming from the academic world, there is often a difficult, and not always pleasant, transition in mindset necessary to be successful in business.

DARPA generated another problem for the AI firms. In the late 1980s DARPA began cutting much of its research dollars which had, by themselves, accounted for a large amount of the artificial intelligence market. While AI firms, both hardware and software, could have subsisted on corporate clients alone, they were shackled by their in-grown dependence on government clients. This dependence stemmed from their academic origins; early on, the firms with stronger links to the "big three" academic AI labs did better (see

Inference Corp. section) in winning the big early deals from the government (which was both a supporter of research and a large customer of AI products). However, these firms later found it harder to make the switch to serving purely corporate customers.

3.2 Business Models: Products Versus Consulting

A major problem for many of the expert systems (ES) software firms (Teknowledge, Carnegie Group, etc.) was their underestimation of the effort required to implement their systems. They thought they could sell their software as a product, and with a minimal amount of training, let the buyer's IT department install and set up the software. However, these pieces of software were terribly complex and required large amounts of customization and knowledge entry, especially if the buyer wanted them to work well. They were designed by and for the best few programmers in America's top computer science departments. Most IT departments had no chance without large amounts of consulting help, which most of the ES firms found themselves providing.

However, this help came at a price to those firms. Consulting firms cannot achieve the same levels of potential profit as product firms, because the revenue can only grow with the number of people that they hire, whereas a software product has almost no marginal cost of production. Venture capitalists like to see product firms in software for exactly this reason, and the venture-backed ES companies found themselves facing pressure from their venture capitalists to produce more products. But if they wanted their companies to succeed, they had to actually deploy, or at least install, some of their software, which required them to do this consulting work.

3.3 Misunderstanding the Target Market

In retrospect, one of the biggest issues the AI firms faced was that they did not have a very good understanding of what their mainstream market was looking for. While they did reasonably well with early adopters of their technology, most of the early firms never changed their strategy to sell to Moore's "pragmatic" market. These marketing problems are explored in more detail below.

3.3.1 Incompatibility with Clients' Internal Systems

The ES companies faced another challenge, which many did not realize until it was too late. They had handcuffed themselves to the Lisp programming language. In the academic world, Lisp was generally considered one of the top languages for working on AI systems, since it uses data tags that keep the data type separate from the data itself. This fact makes it slower to run on conventional architectures, but gives more freedom for the programmer who is writing in it. Lisp thus gives the programmer the flexibility to manipulate both programs and data, and eases the rapid prototyping of software programs.² Even today, MIT teaches its introductory computer science course, as well as its artificial intelligence courses, in a dialect of Lisp.

However, in the corporate world, Lisp is an anomaly. Few large-scale systems are written in the language; few large software firms providing languages (namely Microsoft) provide any kind of Lisp support. Thus selling software written in Lisp in which the department customizing it has to use Lisp is very difficult. IT departments, in general, want to minimize the number of things their people have to know; adding more programming languages to that list is generally not well received. The same is true on the hardware side: most large IT departments try to stick to one (if they just want simplicity) or two (if they want to provoke some competition from their suppliers) major vendors of hardware, and the specialized hardware firms found them thus a tough sell.

Ideally, a piece of expert system software does not sit secluded from the other software systems in the corporation. The corporation's other databases and systems contain much of the data that the expert system should work on, and thus ideally the ES will link into those systems and gather its data that way. There are many benefits to designing the ES this way, as opposed to making it stand alone, not the least of which is that updating the data only has to occur once. However, being written in Lisp meant that the IT department had to build bridges from the Lisp code to the databases and other software that the rest of their systems were written in (which could be FORTRAN, C, PASCAL, etc., but definitely not Lisp). Sometimes the ES firm would provide these bridges, but often IT departments would have internally written software and thus the IT department would have to build its own bridges. This feat, obviously, required the ability to write code in Lisp. And unfortunately, Lisp programmers were not easy to find; and when they were found, they

tended to command very high salaries. The result was that very few of the expert systems built in the 1980s were actually deployed.

Eventually many of these firms switched to Lisp products, and the ones that did so earlier or started out in standard languages (like Neuron Data) did better than those that delayed the change (Teknowledge, Gold Hill). There was a large amount of hubris, not completely unwarranted, by the artificial intelligence community that Lisp would change the way computer systems everywhere ran. Too late they saw that they were the Mohammed and corporate IT departments were the mountain, and the mountain was not going to move to them. Like Copernicus, the AI community needed to realize that they were just one more planet revolving around the sun, not the sun itself. The Copernicus concept was not heeded for another half dozen years, much to the detraction of the AI industry.

3.3.2 Selling Tools Versus Vertical Solutions

An early lesson in economics teaches that when two producers, who have different costs in producing different goods, each produce the good for which its costs are cheapest, the market is at its most efficient point. Looked at another way, this point suggests that companies should focus on providing solutions that they can best provide, not tools for their customers to build their own solutions, especially when there is a large amount of expertise required to build the solution from the tool. Solutions, not tools to let someone provide their own solutions, turned out to be where the money was in the artificial intelligence industry, because the effort for most of their customers to build working AI systems from these toolsets was far beyond their IT capabilities. Tools can be a good market in other industries, where the customer has most of the required skills to put the tool to work and needs to customize the end product to their own particular use. Plus, it is important to note where the value proposition was in their industry; in AI, much of the value was created in that last step of implementing a working system. This value should translate into profits for the company's revenue sheet.

The ES firms in particular were guilty of ignoring this concept. Their expert systems were so general, and thus required so much customizing and knowledge acquisition on the part of the customer, that the systems appealed to a much smaller market, the corporate equivalent of the "do-it-yourselfers". Today, many modern ES firms, like i2 and Trilogy Software, have verticalized their product offerings to a specific market segment. By doing this they are able to encode the knowledge into their product and minimize the

customization required by the customer. However, even these firms have found a large amount of consulting work is required in order to actually deploy systems, and have worked that into their business models.

3.3.3 Selling Technology Versus Real Products

A classic (in the sense that all high-technology firms struggle with this issue) problem that the AI firms each faced was that they were so excited about their technology, they forgot that their customers wanted solutions to their problems. Whether that solution was low-tech or high-tech did not really matter; that the problem was solved quickly, cheaply and easily was critical, however. Many of the executives of these AI firms were coming from the academic environment which was much more technology focused; gaining contracts from DARPA was still more a technology undertaking than a solutions one.

Thus the sales approach of the ES firms, which need not have focused on the technology much at all, spent most of their time talking about AI. Way too much time was spent debating the relative merits of forward chaining versus backward chaining (different techniques for finding solutions in the ES), instead of what problems they were solving for their customers.

3.3.4 Hardware Insufficiency and Cost; Moore's Law

The hardware firms in question faced the same set of issues that any specialized hardware maker faces even today: Moore's law. The generally accepted form of this law states that the number of transistors per square inch on integrated circuits (on general-purpose chips) will double every eighteen months; most take this today to mean that processing speed will double with the number of transistors. A specialized hardware manufacturer expects to find a market based on the fact that a general chip will do many things well, but nothing particularly fast; its designers, in general, purposely do not make tradeoffs in favor of improving performance for some particular function. Specialized hardware manufacturers seize on this opportunity by seeking out markets that crave better speed for some function, and they build chips that perform that function very well, although at the expense of some other functionality. These specialized hardware manufacturers, while facing a smaller market than the general hardware manufacturer, can nevertheless charge much higher prices to their customers who are willing to pay a premium for that improvement. Thus graphics chips, super-fast Cray computers, and, at one time, Lisp machines were all able to carve out a market for themselves.

But once the general-purpose chips improve to the point of matching the performance of the specialized chip, most customers will switch to the generic machines because their prices are so much lower. The specialized hardware manufacturer can try to continue to improve its chips at the same rate, but are often less well capitalized than the generic hardware manufacturers. This battle is often a difficult one to fight, as many of the Lisp hardware manufacturers soon found. Also, without adequate software to be able to connect the specialized machines to the customer's general machines (where much of the customer's important data is often stored) the usefulness of the specialized machine is limited. And, as we shall see, the Lisp software vendors were late in creating products to connect these machines together.

3.3.5 Missing the PC Trend

The AI hardware industry, as well as several of the software firms, also suffered from missing one of the key turns in the advance of computer market: the rise of the PC. While a few software firms jumped on the PC bandwagon early on, such as Gold Hill Computers with AI tools and Neuron Data with expert systems shells, many of the software firms stuck with the Lisp hardware makers. As consumers and businesses move towards Intel-compatible PCs and Microsoft operating systems, the Lisp hardware and software platforms became less and less palatable. Especially as the performance of AI software on PCs began to compare to that on the specialized Lisp chips, there was little reason for customers to commit to a new platform.

3.4 Failing to Manage Expectations

In some respects, the press created and then destroyed the artificial intelligence industry. The press had enjoyed writing about the prospect of intelligence machines for decades, especially since the release of *2001: A Space Odyssey*, and Time Magazine picking the Computer as the Machine of the Year in 1983. With the publishing of Ed Feigenbaum's book on the Fifth Generation, suddenly they had something they thought was real to write about. The early successes, such as Stanford's MYCIN (to aid physicians in selecting antibiotics for their patients) and DEC's XCON, helped add to the buzz.³

The executives of the AI firms were not about to slow down the hype; this hype was bringing customers to their doorsteps and funding into their coffers. But the hype itself was always uncertain; AI was either booming or dying, seeming to bounce back and forth every year or so. In 1985, writers were warning against reliving the intelligent-machine hype of the 1950s;⁴ later that year another article claimed

AI was out of favor for venture capital investing.⁵ Then later that year at least one writer (from one of the same papers that was calling for its demise) claimed AI would be one of the "most likely fast-growth areas."⁶ In 1986 things were definitely hot again, as "The Gang of Four" expert systems companies (IntelliCorp, Teknowledge, Carnegie Group and Inference) was showered with more publicity,⁷ proclaiming "artificial intelligence is hot"⁸. In 1986 projections for the 1990 AI market ranged from \$2 billion (from Financial World Magazine) to \$12 billion (from Arthur D. Little). Actual revenue numbers averaged around \$400 million.⁹

But the hype did not stop anyone; never mind that they could not actually deliver what these customers expected. But through articles in the early 1980s in Fortune, Forbes, The Wall Street Journal, even the conservative world of finance became interested. Several firms managed to have IPOs in the late 1980s, taking advantage of the hype and excitement (it certainly was not their revenue numbers that led to successful initial public offerings). Those that did not get out then found that a few years later, the market had already cooled to artificial intelligence. Once the market started turning south, so did the press: those same newsmagazines that were singing the praises of AI a few years before, were now drying tears at its funeral.

In the end, this hype propelled the AI industry past its problems for the first decade. Customers, both corporate and government, bought AI hardware and software on its hype and its promise. By the time the industry's problems, both those internal to the companies and those external in its market, caught up with it, the press had turned sour and it was too late for them to save themselves in their current form. Thus some companies disappeared altogether; others struggled along, trying to find their niche, and even exist today, often in some diminished or greatly altered form.

¹ Sources: Newquist, *The Brain Makers*; Mark Clifford, "Artificial Intelligence Investing in High Tech Firms," *FW* (23 January 1985), 13; Emily Smith, "A High-Tech Market that's Not Feeling the Pinch - Eager Investors Have Created a Boom in Artificial Intelligence," *Business Week* (1 July 1985), 78.

² Peter Norvig, *Paradigms of Artificial Intelligence Programming: Case Studies in Common Lisp* (<http://www.norvig.com/paip-preface.html#whylist>).

³ Richard O. Duda and Edward Shortliffe, "Expert Systems Research," *Science* (15 April 1983), 261.

⁴ Mark Clifford, "Artificial Intelligence (Investing in High-Tech Firms)," *FW* (23 January 1985), 13.

⁵ John Eckhouse, "Hot Investments for 1985," *The San Francisco Chronicle* (1 February 1985), 35.

⁶ Daniel Rosenheim, "Silicon Valley Slump -- It's Not All That Bad," *The San Francisco Chronicle* (23 August 1985), 6.

⁷ Matt Rothman and Emily Smith, "The Leading Edge of 'White-Collar Robotics' - These Hot Startups Are Rushing to Cash In On Computer Software That Mimics Human Reasoning," *Business Week*, (10 February 1986), 94.

⁸ William Bulkeley, "Stocks of Artificial Intelligence Firms Prosper, Though Some Analysts are Advising Wariness," *The Wall Street Journal* (31 March 1986).

⁹ William Bulkeley, "Bright Outlook for Artificial Intelligence Yields to Slow Growth and Big Cutbacks," *The Wall Street Journal* (5 July 1990), B1.

4 The AI Software Industry

The 1980s AI software companies that I am focusing on were producing products in one of the following areas: AI programming languages and tools; natural language software tools, expert system shells; and application-specific customized expert systems. Other types of artificial intelligence applications, such as those related to vision systems, robotics, and neural networks, also existed but did not have the same levels of activity as the areas I have chosen.

Many of these firms decided to sell various software tools that would allow their clients to build their own systems. Most soon learned, however, that selling tools is a difficult business model, and these corporate customers were not prepared to do their own development. A software company can get a very high price on its technology in an application, but it must be packaged. Palladian, one of the expert system application firms, understood this part of the model, but still stumbled, possibly due to their use of academic, instead of industry, knowledge. These firms should have been hiding the technology from the customer, and codifying the non-customer-specific knowledge within the program with most of the AI decisions already made; business people should have been able to provide their particular information without understanding the technology. Unfortunately it took a painful decade for the industry to get to that point.

4.1 Background of Expert Systems

Expert systems are pieces of software that include, generally, a database of facts; a set of rules; and a way for users to enter the specific data of their problem. Once the user inputs their specifics, the software applies the set of rules and its own knowledge to output an answer to the user's question. For example, a financial planner may input the salary, net worth, and risk profile of his client, and the system, having rules for when to invest in different financial instruments and data about various particular securities, could output a set of securities for the client to invest in.

The companies providing software in the expert system space generally fall into one of two areas: expert system shells and expert system applications. Shell companies, such as the "Gang of Four," mainly offer a structure for knowledge representation and an inference engine, but the user must supply the

knowledge. These systems usually required large amounts of work in encoding the knowledge and implementing the system, either from the client's IT department or from the AI firm's consulting unit. Early on, shells were commonly written in the Lisp programming language, which was considered very good for AI application but not used much outside academia; Prolog, OPS, and C were also used. Later on, C++ also became popular.

Expert system application firms verticalized their offerings into a specific area, such as the financial planning example above. Firms like Palladian encoded knowledge in a particular area and sold the system mostly "as-is," although clients would often want to connect the ES into their corporate databases.

The common problems that expert systems attempted to solve were configuration (such as finding possible configurations of DEC's VAX machine); scheduling (such as planning various tasks within a large project); classification (of, for example, various chemical compounds); interpretation (looking at a series of events and determining their meaning within a certain domain); and diagnosis (by, for instance, looking at symptoms of a disease and interpreting the meaning).

Table 4.1 summarizes the above information and describes the building blocks of expert systems, from low-level programming languages and tools to high-level applications.

Table 4.1: Expert Systems¹

Product Area	Examples	
Domain-Specific, Task Specific Applications	XCON; Authorizer's Assistant, MYCIN	High-level tasks
Tasks	Configuration; Scheduling; Classification; Interpretation; Diagnosis	
ES Shells	S.1, M.1; KEE; EMYCIN; ART; Nexpert; KnowledgeCraft	Low-level tasks
Languages	Lisp; OPS; Prolog; C	

In 1985, the expert systems market reached \$74 million in sales; the next year they projected 1990 revenues at \$800 million.²

4.1.1 Early Expert Systems

Several early expert systems were instrumental in starting the hype in first the academic community and then the corporate world. This section describes a few of these systems.

MYCIN

MYCIN was developed at Stanford University in the 1970s as a physician's aid to select antibiotics for their patients. Knowledge of infectious diseases is encoded in this rule-based system, and based on input from the physician on the results of various tests. For example, the system might determine "if (i) the infection is meningitis and (ii) organisms were not seen in the stain of the culture and (iii) the type of infection may be bacterial and (iv) the patient has been seriously burned, then there is suggestive evidence that *Pseudomonas aeruginosa* is one of the organisms that might be causing the infection."³ Mycin's accuracy in recommending antibiotics was comparable to those of a physician.

DENDRAL

Another expert system produced at Stanford in the 1970s, DENDRAL analyzed mass spectral patterns to determine the compound's chemical structure. It works on multiple families of compounds and contributed to several journal articles.⁴

PROSPECTOR

PROSPECTOR was a mineral exploration system used for evaluating resources, ore deposit identification, and selecting drilling sites. Its knowledge base contains models of ore deposits. Its performance was quite close to that of geological consultants.⁵

4.1.2 Issues in Building Expert Systems⁶

Knowledge Acquisition

Although most early expert systems were based in relatively narrow domains, even ensuring that the information that they had was complete was a challenge. It was always a challenge to find experts who could express their knowledge in a way that the programmer could encode into the system. There is also the issue of changing and updating information, either through some sort of learning system or directly from those that upkeep the system.

Knowledge Representation

In most expert systems, knowledge is encoded in rules, although the artificial intelligence community is quite interested in this problem. The main conflict is simplified as being between complex representations that effectively reflect each individual situation and more general representations that is easy for the programmer to interpret and amplify.

Inference and Uncertainty

Most systems must employ some technique for dealing with situations when not all of the input data is available. They thus employ heuristics to “guess” at which answer is best. Designers have their choice of various techniques, such as Bayes’s nets, possibility theory and Dempster/Shafer theory of evidence.

Explanation and Interface

Early on, expert system designers discovered that systems that gave results without explaining their reasoning caused their users to distrust the results. Thus systems need to find ways to track its reasoning and explain their path. Beyond the system’s explanation is also the usability challenge of making the interface simple enough for users who may not be familiar with the technology to input the specifics of their situation and understand the results. Designers must also make the data and the rules of the system reasonably easy to update and correct.

4.2 The Gang of Four

The first group of expert system shell companies earned themselves the dubious nickname of the "Gang of Four" after the clique of

radical advocates of Mao Zedong who implemented the most extreme policies of China's Cultural Revolution during the 1960s and 1970s. The group consisted of Jiang Qing (Mao's third wife), Wang Hongwen, Zhang Chunqiao, and Yao Wenyuan. All held only inconsequential political power prior to 1966 when the Cultural Revolution began. Zhang and Yao were minor propaganda officials in Shanghai. ... The members of the Gang of Four emerged as Mao's principal supporters in the campaign and were rewarded with increased power. By 1969 all were members of the ruling Politburo of the Chinese Communist Party (CCP). Jiang was especially valuable to Mao as a trustworthy ally against the moderates.

The Gang of Four first began to act collectively in 1965 when Yao published an attack on a play by Wu Han that Jiang was investigating for promoting counterrevolutionary ideas. The incident was one of the triggers for the Cultural Revolution. ... At the Tenth Party Congress in 1973, Wang emerged as heir apparent to Mao and first premier Zhou Enlai. Mao's death on September 9, 1976, however, removed the Gang's main source of power. They were arrested and charged with various crimes, including treason

and forgery of Mao's instructions. Cartoons and other attacks vilifying them spread in the media and the term "Gang of Four" was adopted for them.⁷

All four were imprisoned for life. The treatment for the AI Gang of Four, which consisted of the Carnegie Group, IntelliCorp, Teknowledge and Inference, was somewhat better; all four firms still exist in some form today, either independently or as a division of a larger firm. However, none had the skyrocketing success that the hype around them predicted.

Table 4.2: Gang of Four Snapshot, 1986⁸

<i>Firm</i>	<i>Founded</i>	<i>Funding (as of 1986)</i>	<i>Product</i>	<i>Major Investors</i>
Carnegie Group	1983	\$11 million	Knowledge Craft	DEC, TI, Boeing
IntelliCorp	1980	\$29 million	KEE	Public
Teknowledge	1981	\$17 million	S.1	GM, P&G, Nynex
Inference	1979	\$15 million	ART	Lockheed, Ford

4.2.1 Carnegie Group Inc.

CMU's Turn: Why let MIT and Stanford Have All the Fun?

The Carnegie Group, as the name suggests, was spun out of Carnegie Mellon University in 1983 by four CMU scientists, Raj Reddy, Jaime Carbonell, Mark Fox, and John McDermott to commercialize the AI and natural language technologies they had been developing. Realizing they needed more professional management, they chose entrepreneur Larry Geisel to be its President and Chief Executive Officer (CEO). They decided to focus on Lisp-based expert systems for industrial and manufacturing use. By 1985 the press called Carnegie Group "Pittsburgh's premier artificial intelligence firm."⁹

Playing With the Big Players

More than any other major AI firm, Carnegie took to selling pieces of its equity to large corporations in return for cash and product testing sites. The first big investor was Digital Equipment Corp, who, after XCON, saw Carnegie as a way to further its own AI research without the bureaucracy that hindered its in-

house efforts. It paid \$2 million for 10% of Carnegie in 1984.¹⁰ Later investors included Ford Motor Company (\$6.5 million), US West Inc. (undisclosed amount), Boeing Co. (\$1.6 million), and Texas Instruments (\$5 million);¹¹ by 1991 outside investors owned 55% of Carnegie's total shares.¹² Being this dependent on its larger partners, however, made it difficult for the firm to focus on developing generic products for the marketplace; most of the work with these firms was done on a client-specific basis.

Carnegie did a few initial products, which were written in Lisp. These included Knowledge Craft, a set of tools for creating large expert systems, and Language Craft, a software environment for creating natural language interfaces to other applications, databases, and knowledge-based systems. By 1986, Knowledge Craft was available on DEC's AI-based VAXstation, and HP's Series 9000 Model, and Language Craft was available on Symbolics, VAX, and TI Explorer systems.¹³

The First Reckoning

In 1987, Carnegie Group's "hand-picked" President, Larry Geisel, resigned¹⁴ to found another AI firm,¹⁵ Intelligent Technology Group, which sold AI-enhanced software for investment portfolio management to the banking industry. Unfortunately for Geisel, ITG filed for bankruptcy protection in 1991 when a large enough market for its products failed to materialize.¹⁶ One of the technical founders, Mark Fox, took over Carnegie Group, and later that year he brought in Dennis Yablonsky, former President of Cincom Systems, to fulfill the President and CEO slots. Yablonsky brought a much-needed sales and marketing oriented approach to Carnegie.¹⁷

In 1987, Carnegie bought "The Operations Advisor" for \$30,000 from Palladian, a troubled AI software company that was having troubles staying afloat. But "Carnegie couldn't make [Palladian's software] work"¹⁸ in its current form. They ported the product to the PC and renamed it Operations Planner, and they dropped the price to \$4,000.¹⁹

Carnegie Group continued to build successful custom expert systems for its partners; unfortunately, this work did not necessarily lead to profits for the firm itself. In 1990 they even had a Harvard Business School case published on them. The year was 1989, and Carnegie had never posted a profit. The case asked whether Carnegie should attempt to start an initiative, called CORE, to build a commercial technology with all of its major partners, instead of continuing to work on the bilateral projects (wherein Carnegie teamed with just one of its partners) they had done in the past. Their relationships with their

various partners differed greatly: DEC was interested in advanced tools; the rest mostly wanted applications. US West and Ford wanted to use the technologies, while DEC and TI wanted to sell the products.²⁰ These variations made it difficult for Carnegie to focus its product strategy.

Nevertheless, the result of the CORE initiative was that in 1990, Carnegie announced the formation of the Initiative for Managing Knowledge Assets (IMKA) with DEC, Ford, TI, and US West. Their goal was to develop a new knowledge based system technology. By this time, custom software applications were still producing 70% of Carnegie's revenues, although that was down from 90% in 1986.²¹

Under Yablonsky, Carnegie's fortunes improved. In 1992, Yablonsky was named Entrepreneur of the Year in the Turnaround Reorganization Category by Merrill Lynch and Ernst & Young, Inc. magazine. They had grown to 175 employees and achieved eight quarters of profits.²² Shortly thereafter, Carnegie signed a deal with Caterpillar to develop a machine translation system for Caterpillar's technical documentation,²³ and the firm announced a new release of ROCK, the result from the IMKA association, a product for processing and storing complex and dynamic information.²⁴

Basking from the glow from these positive steps, Carnegie debated taking that next step and going public. The management waited until 1995 to take that step, and at \$8 a share, it raised \$11.1 million in its IPO.²⁵ Despite the cash inflow, Carnegie continued to have problems with its dependence on key clients. In 1992, Carnegie lost its defense business, which was 40% of its revenue, and the firm laid off 20% of its employees. By 1994 it was growing again, but in 1996, a major telecommunications client cancelled its contract and took with it 30% of Carnegie's revenue.²⁶ Carnegie attempted in 1997 to refocus on customer interaction and logistics, planning and scheduling, but they did not have the resources to fully implement this strategy, and its stock continued to languish, dropping as low as \$1.75 per share.²⁷

Carnegie's days as an independent company finally ended in 1998 when Carnegie was purchased by Logica plc for \$35 million. By this time, Carnegie had grown to 300 people and specialized in customer relationship management software and decision support solutions.²⁸

4.2.2 IntelliCorp Inc.

Founding: IntelliGenetics

In the late 1970s, the Stanford University Computer Science department built several expert systems in conjunction with the medical school, namely DENDRAL, MYCIN, and, starting in 1975, MOLGEN.

MOLGEN was built to help researchers analyze DNA sequences. This program was very successful, in that researchers linking in to the computer to run it began to strain the department's resources. In 1980 Stanford medical school researchers Laurence Kedes, Douglas Brutlag, and Peter Friedland joined with the computer science department's Ed Feigenbaum (of Fifth Generation fame) and decided that there was a market in putting genetic engineering software on time-sharing mainframe computers for researchers to dial into and use.²⁹

The four men thus founded IntelliGenetics in September 1980 just down the street from Stanford University. They sold licenses to their expert systems cloning software. However, the research market was small, and they began to look for other markets. They also recognized that selling this service, namely access to the software, did not have the same revenue or profit potential that a product-based business did. They considered selling their cloning software as a bundle with a Sun workstation, but the difficulties involved in being both a hardware and software company seemed too great to follow this path.³⁰

Transition to Expert System Shells

Feigenbaum had been involved in the creation of MYCIN, and they recognized that a stripped down version, EMYCIN, which was the expert system shell without the knowledge base, could be used as a product with any knowledge based plugged in. They built a shell, and called it the Knowledge Engineering Environment (KEE). KEE was written in Lisp and ran on the Symbolics 3600 and the Xerox Dorado 1108 machines; it was introduced at the end of 1983.

In 1983, IntelliGenetics also sold a chunk of equity for \$1 million to Computer Services Corp (CSK) of Japan. This agreement gave CSK the right to sell IntelliGenetics' software in Japan. At the end of the year, however, the firm needed to raise more money. Much like many of today's Internet companies, IntelliGenetics had a successful initial public offering (IPO) on its "sizzle" in December 1983, raising \$9 million. IntelliGenetics was the first "AI" company to go public, as a company that "makes computer programs based on artificial intelligence technology for biotechnology and other applications."³¹ In the summer of 1984, cementing its movement towards generic expert systems shells, IntelliGenetics changed its name to IntelliCorp.

On the management side, IntelliCorp brought in several people from Texas Instruments, including Gene Kromer and Tom Kehler. Kromer became President in 1984 and replaced Tony Slocum, who left to

form Lucid, and Kehler, who originally oversaw the firm's business development activities, and later became CEO,³² brought Greg Clemenson with him from TI for the technical side. Clemenson was instrumental in building their shell; his focus was on the knowledge representation framework, and used a simple rule engine (other firms, such as Teknowledge, focused more on the rule engine).³³ Eventually, they gave up on the biological software side of the business entirely, and sold 60% of the IntelliGenetics subsidiary to Amoco in 1986, and sold the remaining interest to them in 1990.³⁴ Four years later, this subdivision was bought by the Oxford Molecular Group.³⁵

Building a Market

IntelliCorp attempted to train its customers in building AI applications through its "apprenticeship programs", but it turned out to be much harder to pass on the skills of expert system building than they had supposed. Most of the successful products that were built from KEE were either built by IntelliCorp engineers or with their close help.³⁶ This training was reasonably lucrative, however, as they were able to charge three times the industry average for their training services.³⁷

In 1985, IntelliCorp signed a licensing deal with Sperry that allowed them to market KEE in return for \$4 million and consulting work.³⁸ Sperry was working with Northwest Orient on the SeatAdvisor system, which was built using KEE and ran on TI's Explorers. This system helped the airline extract the largest possible price per seat on a continuous basis. However, in the middle of production, Northwest Orient merged with Republic Airlines, and the SeatAdvisor project was deemed too distracting and stopped, to the chagrin of TI and IntelliCorp.³⁹

IntelliCorp ported KEE to the Sun platform, but keeping track of the various version numbers was a full time job, thanks to the many different flavors of Unix. It took a very nimble firm to stay on top of all of these nuances, but IntelliCorp did.⁴⁰

IntelliCorp's public status was useful in gaining notice; in November 1985, it was called a high potential stock in the field of artificial intelligence by the Chicago Sun-Times despite posting no earnings since 1980.⁴¹ Later that year, it had a second offering that raised \$22.7 million, this time with top-rated investment bank Montgomery Securities. In 1986, IntelliCorp began moving towards more generic hardware architectures by releasing KEE PC-Host, which enabled its customers to run its programs from PCs connected to mainframes, though the programs still had to be written in Lisp.⁴² The market seemed to

applaud the move; that year they traded at about 50 times projected earnings.⁴³ In 1987 IntelliCorp announced that KEE could connect to mainstream databases.⁴⁴ And the next year, the company announced that KEE itself would be available on IBM-compatible PCs, albeit very powerful ones (10 megabytes of memory, 100 megabytes of disk space) for the era.⁴⁵

Endgame: Shifting Away From Expert Systems

In 1989, IntelliCorp earned almost \$1 million in profit on revenues of \$22 million. However, this was its last profitable year. In 1990, it acquired MegaKnowledge and its KAPPA object oriented tool, as part of its strategy to de-emphasize KEE. This new direction was not taken well by the Lisp group (KEE was one of the last expert systems still written in Lisp, after most of the rest had switched to C). This transition also confused IntelliCorp's customers, who stopped buying. The CEO, Tom Kehler, decided to sell the company, and set up a deal with KnowledgeWare to sell it for \$34 million in August 1991. However, in November, KnowledgeWare announced an unexpected quarterly loss, and the deal fell through. IntelliCorp's board was furious, and Kehler left the CEO position; COO K.C. Branscomb took over.

However, she was unable to save the company, and she left in October 1992 (although she stayed on as a director) and the firm was taken over by its CFO, Kenneth Hass.⁴⁶ It refocused on KAPPA, and today, under CEO Haas, the firm develops enterprise resource planning software.⁴⁷

4.2.3 Inference Inc.

Inference was founded near Los Angeles in El Segundo, CA, by Alex Jacobson and Chuck Williams in 1979, making them the first of the original "Gang of Four" expert systems companies.⁴⁸ Unlike the others, Inference had no strong ties to the academic "AI Mafia" of researchers from MIT, Stanford and CMU; Williams, its CTO, held a bachelors degree in computer science from the University of Southern California and had conducted AI research at the USC/Information Sciences Institute. Also, the firm was located in southern California, not Boston or Silicon Valley, where most of the high-tech startups were found. In the early days, these facts often caused them to be discounted by the growing AI industry; they also found it difficult to obtain DARPA work, although they did do some. However, this fate ultimately had the benefit of decreasing their dependence on DARPA and forcing them to find more private companies for whom to build systems, which helped them survive when DARPA funds dried up, causing many of the other expert systems companies to falter.⁴⁹

Focus on Applications

From the start, Inference strove to build applications, not just toolsets. In 1983, Williams was quoted in the press saying that a development tools strategy would not work (Inference focused on applications). However, the government's sponsorship and involvement in the industry induced the market for tools and created an artificial market that drove the early success of all four ES companies.⁵⁰

In the early days of artificial intelligence commercialization, the DARPA Strategic Computing Initiative (SCI), which began at least partly in response to Japan's Fifth Generation Project, gave \$700 million to aerospace firms to become expert in AI. Inference, however, decided very early that they did not want to live on just the aerospace market. They would try to sell to commercial accounts, like American Express, which resulted in the successful Authorizer's Assistant program, which helped Amex determine whether credit-card purchases should be approved.⁵¹ Through work with industry, they learned they needed a different infrastructure to build their technology on, namely generic hardware (mainframes) and software (C), not the Lisp and Symbolics machines that most of their competitors used.⁵²

Funding

For funding, Inference, like many other AI firms, formed agreements with larger industrial companies. In 1984 Lockheed assumed a minority interest in Inference,⁵³ and in 1986 put \$2 million more in, raising its investment total to \$6 million.⁵⁴ In 1985 Ford put about \$14 million into Inference (as well as a similar amount to the Carnegie Group) in the form of equity, development contracts in financial services and industrial engineering applications, and technology transfer agreements. Inference was to build several expert systems, one for approving credit, another for the design and diagnosis of brake systems, and a third for industrial engineering stands in the manufacturing process.⁵⁵

Inference did not just use their corporate investors; they also took around \$30 million of venture capital funding, in at least eight rounds.⁵⁶ Their venture capital investors included JP Morgan Capital, Venrock Associates, and Corporate Venture Partners, as well as Lockheed and Ford.⁵⁷ The venture capitalists, at various stages, forced management changes on the company, some with bad results; the new management would then change the company's direction. For example, the 1991 executive that the venture capitalists brought in changed the direction of the company positively towards the client/server direction, and he also took the company public.⁵⁸

Products and Customers

Inference's first product, introduced in 1984, was an expert system shell, the Automated Reasoning Tool (ART). Originally written in Lisp for Symbolics workstations, it competed with KEE from IntelliCorp.⁵⁹ Their big coup came in 1986, when they won a deal with American Express to build their Authorizer's Assistant program, which ran on a Symbolics Lisp machine. This program required a fair amount of consulting from Inference to complete, and the first prototype, with 520 rules, took six months.⁶⁰ The productivity savings alone from Authorizer's Assistant generated a 45-67% internal rate of return (IRR) for American Express.⁶¹

Despite their corporate successes, they continued to have problems in the government sector. An example of this difficulty presented itself in their agreement with NASA. Inference's first customer was the NASA Johnson Space Center, which was building the space shuttle control function for ascent and descent to control the shuttle around the globe using ART. NASA wanted to implement the system in the Mission Control Center. But ART ran on a Lisp machine and NASA needed a PC version, since NASA's systems were highly regulated. NASA built a clone of ART called CLIPS and began licensing it at fairly inexpensive terms to industry. Today CLIPS is widespread; Calico, for example, uses CLIPS technology.⁶²

In early 1987, many of the AI firms started getting into trouble when government interest in funding their projects waned.⁶³ Inference survived, but not because it never took on government projects. Their successful systems included DARPA's Pilot's Associate project for fighter aircraft pilots, and two systems at the Air Force, one for ensuring the availability of trained personnel for various needs, and another for handling the availability of F-16 jets.⁶⁴ Inference survived by porting ART to almost every available computer architecture; this strategy can be risky if the porting effort sacrifices new product development, but Inference managed both. By 1989, ART was available on DOS, IBM MVS, IBM AS/400, and DEC VAX environments.⁶⁵

Using Consulting to Stay Afloat

The consulting unit was made a formal business unit in 1988. The unit kept the company alive for several years, from 1988-1990; it was profitable when nothing else was. In fact, the consulting unit provided the bulk of the revenues and all of the profit during this time period.⁶⁶ In 1990, within the consulting unit, a research unit started looking at new technology, called Case Based Reasoning. They saw

lots of interest for this technology in the customer service area, and they decided to start a "skunkworks" project to build a prototype of a software system for customer service representatives.⁶⁷

In 1991, Inference brought in commercial software management and made another play at the tools market. Seeing the client/server revolution coming, they recognized a need for new class of development tools for client/server systems with friendly user interfaces, database access and AI capabilities. They tried to build a tool as powerful as ART (which they found they could do) and as easy to use as PowerBuilder (which turned out to be very hard); unfortunately the product did not take off in the marketplace.⁶⁸

The Rest of the Story: Call Centers and Brightware

The Call Center application, DBRexpress, which helped companies manage their corporate call centers and respond to customer inquiries, took off in the market. It was easy to implement and maintain, and there was a clear market to which to sell it. In 1995, the part of the company with the call center application (which was called Inference) went public. Williams (then the CTO) spun out the rest of the firm, which included the most recent version of ART, into a separate firm that would build a new set of tools. This firm, called Brightware, now develops software for email customer service, and still sells ART*Enterprise for custom application development.⁶⁹

4.2.4 Teknowledge Inc.

After the founding of IntelliGenetics, several Stanford researchers decided there was more money to be made in selling expert systems shells. In 1981, twenty of them (including Ed Feigenbaum and Peter Friedland from IntelliGenetics; Randy Davis, who would go on to found Applied Expert Systems; Jerrold Kaplan, later of GO Corp, author of *Startup* and currently running OnSale; Douglas Lenat, later of Cyc fame; and Frederick Hayes-Roth, who had worked on CMU's HERESAY-II project) founded Teknowledge to sell knowledge engineering services. Realizing they needed more professional management, they hired Lee Hecht, a former university lecturer and founder of several cash management and one motion-picture company, to be CEO. Their first year was spent mainly doing consulting work.⁷⁰

But the company realized it needed a product if it wanted to be a successful software firm, and in 1984 announced their first product, an expert system shell for the PC called M.1. Compared to IntelliGenetic's KEE, M.1. was a low-end product; it cost \$12 thousand, which was quite reasonable, one-fifth the price of KEE's \$60 thousand price tag.⁷¹

They hired several salespeople straight from business school to sell M.1 to corporations.⁷² In an classic example of their mistakes due to their newness to business, the salesmen, thanks to a lack of time to develop a proper demonstration, used a program called "The Wine Advisor" to display the product's features. This program took as input various meal options and then chose a wine that would best suit that meal. However, this demonstration failed to impress the prospective corporate clients, and made the product, at least initially, rather difficult to sell. However, thanks to its relative cheapness (compared to KEE) and PC platform, it did begin to sell.⁷³

About this time, Teknowledge began looking for investment, and in 1984 sold 11% of the firm to General Motors for \$3 million,⁷⁴ later raised to \$4.1 million. In time, Procter and Gamble put \$4 million into the company, NYNEX put in \$3 million, and FMC Corp put in \$3.2 million.⁷⁵ In March 1986, Teknowledge followed the route of IntelliCorp and Symbolics and went public at \$13 a share, 81 times pro forma annualized operating earnings of 16 cents per share.^{76,77}

To expand its product line into the high-end systems, Teknowledge developed S.1, which ran on workstations, and was more powerful and more expensive. Unfortunately, it was not compatible with M.1, meaning customers would have to start all over again to re-code their systems into the more powerful program.⁷⁸

At the end of 1985, Teknowledge realized that the future was in C, and announced it would stop supporting Lisp and PROLOG and would do their work in C. They were the first major AI expert systems firm to make this change, and it upset much of the rest of the AI community.⁷⁹ However, while it was a step in the right direction, it did not save Teknowledge's AI business. The firm was spending much of its time consulting, not building products; for the last six months of 1985, half their sales were from services, and only a quarter from selling software.⁸⁰

To add to their problems, by 1988, inexpensive expert systems, such as Paperback Software's \$99 VP Expert, were eating away at Teknowledge's market share. Looking at Paperback's products, they saw enough similarities to fight back by suing them. Paperback was also being sued by Lotus, and eventually shut down. Also, the other expert system vendors had made their products run on workstations, thus also hurting S.1.⁸¹

Teknowledge began laying off employees and brought in a new CEO from FMC, Peter Weber. They sold off one of the divisions, the Federal Systems Division, for \$1.5 million to ISX Systems. They began to refocus on providing consulting services, as their products were not selling, but still lost \$10 million in 1988 on \$14 million in revenue. Around the end of 1988, Teknowledge was merged with American Cimflex, which sold computer systems for manufacturing, and had its own balance sheet problems. The new company, Cimflex-Teknowledge, continued to lose money.⁸² In 1993, Hayes-Roth took over as Chairman and CEO of the merged company, returned the firm to the name Teknowledge, and brought the company to profitability by the end of 1994. The company today sells software based on distributed knowledge processing.⁸³

4.3 Expert Systems Shells

The expert system shell business was larger than just the "Gang of Four." Two other interesting firms were AICorp, which was in fact the first artificial intelligence company, and Neuron Data, which started out building expert systems for the Apple Macintosh. AICorp moved to expert systems after a start in natural language processing. Ultimately there were dozens of expert systems companies, including Aion, Advanced Information and Decision Systems, and Knowledge Garden.

4.3.1 Artificial Intelligence Corporation

The Beginning of an Era

Before the Gang of Four, or the Fifth Generation Project, there was the Artificial Intelligence Corporation. Founded in Waltham, MA, by Larry Harris, a former computer science professor from Dartmouth College, in 1975, AICorp began to develop a natural language software product for mainframe computers called Intellect, which hit the market in 1980.⁸⁴

Intellect connected the knowledge inside corporate databases with non-technical users by allowing them to query the database using conversational English. It cost around \$50,000; more than 200 copies were sold by 1984.⁸⁵ That year, they were joined in their field by competitors Mathematica Products, Frey Associates, and Symantec; nevertheless, in June, they signed IBM as a marketing partner. The fact that natural-language programs were rather processor-intensive made them very attractive to hardware companies as a way to encourage more sales of their machines.⁸⁶

AICorp originally planned to develop a PC version of Intellect for release in 1985, but a slumping software market caused them to halt their plans.⁸⁷ The firm did finally turn a profit that year, on revenues of \$4.5 million. By this time, however, the firm had received \$9 million in venture capital, and went through two Presidents until Harris assumed the title.⁸⁸ In 1986, they hired the former President of Cullinet, Robert Goldman, as the new CEO.⁸⁹

Joining the Expert System Shell Space

The company, looking to move beyond just natural language database interfaces, began to ship an expert system shell, called KBMS, in the summer of 1988. They developed KBMS through a consortium set up with Transamerica Insurance Services, Southern California Edison, Liberty Mutual Insurance, and E.I. Dupont de Nemours & Co, all of whom provided both funding and feedback. This product competed with the Gang of Four products, but had the advantage of learning from their mistakes; KBMS included strong ties with other programs using the same data and compatibility within the IBM environment, as well as natural language capabilities.⁹⁰

Acquisitions, IPO, Merging and Renaming

Looking to expand its expert systems product line, AICorp acquired 1st Class Expert Systems in the spring 1990. 1st Class offered expert systems for the low-end PC market, which complimented AICorp's offerings for mainframes at the high end. Together the two companies employed 140 people.⁹¹ In June of 1990, AICorp raised \$21.4 million in a public offering that valued the company at \$80 million.⁹²

In 1992, AICorp merged with Aion, the seller of the direct competitor Aion Development System, and renamed the combined company Trinzic. The new firm sought to include Intellect, KBMS, which targeted developers who wanted to use high level languages, and AionDS, which was aimed at more sophisticated expert system developers, into a broad product suite.⁹³ The combined firm

will continue to develop, sell and maintain a broad suite of product solutions, including its two industry-leading knowledge base system products, KBMS and the Aion Development System, as well as Intellect, a natural language data access tool. Development of AionDS, which is often used by developers who prefer the control and flexibility of detailed programming constructs, will continue to be oriented to the solution of BPA [business process automation] problems. KBMS, which appeals to application developers who prefer to build applications using the power and productivity of high level languages, will continue its migration over time to the object-oriented 4GL marketplace.⁹⁴

Trinzic merged with Platinum Technology, a data warehousing software firm, in 1995. By this time Trinzic itself had moved into the data warehousing space although it was still offering the Aion Development System. After the merger, Trinzic became a Platinum subsidiary.⁹⁵ Today their technology has been effectively merged into the Platinum brand offerings.

4.3.2 Neuron Data, Inc.

AI on a Mac? Ha!

Interest in AI technologies was by no means limited to the United States; however, many abroad saw the U.S. as the market of choice for their products. In 1984, two Frenchmen, Patrick Perez, an MBA who had been on the team to launch the Macintosh in Europe, joined with AI programmers Jean-Marie Chauvet and Alain Rappaport, who had built an expert system in Lisp on a Symbolics workstation to more closely model human problem-solving techniques. They started a company in Palo Alto called Neuron Data.⁹⁶

Putting AI on the Mac

The two saw a large amount of potential in porting the capabilities of the Symbolics machine to the Macintosh, which had similar, albeit more limited, graphics capabilities as the Symbolics machines. They developed a prototype, which they showed to many people at Apple who were quite excited by its potential. However, Steve Jobs nixed the idea, and they went looking for venture capital funding.⁹⁷

They thus developed the first expert system shell for the Macintosh, called Nexpert, written in C (C++ still did not exist, and C allowed them to port to DOS), and shipped in 1986. Their approach to expert systems was generative, solving problems from first principles, as opposed to the rule-based, pattern matching expert systems that most firms used. They wanted to port Lisp to the Macintosh and still keep the graphical views of the Symbolics machine. The firm ended up building an AI interpreter in C that was basically a C executable running a Lisp interpreter. They faced the challenge, however, that they could not keep the functionality or flexibility in a Lisp-based ES of dropping down to the Lisp code since they were writing the product in C; however, Nexpert was still able to keep the graphics and the interpreter features in the product.⁹⁸

The traditional ES firms laughed, as the Mac was looked at as little better than a toy, and few in corporate America used Macs. But in the end, by porting their product to Unix workstations and IBM-PCs

as well as Macintoshes, they were able to sell thousands of copies, and they were one of the few companies to make it through the AI wreckage in the late 1980s.⁹⁹

Porting Nexpert Everywhere

In 1987, DEC started selling Neuron Data's Nexpert tool with its mainframes.¹⁰⁰ In 1988, ND introduced Nexpert Object, an ES product with bridge to Oracle, Sybase and Relational Technology databases for \$5000 for a PC (Mac); it was also available on VAX, Sun, Apollo, and HP workstations. A press release promises that this C-based shell was "an open architecture, unlike Lisp artificial intelligence systems."¹⁰¹ Versions for DOS and OS/2 were available by 1989.¹⁰²

In 1990, they finally ported the product to Windows with the release of version 2.0 of Nexpert Object¹⁰³, and Windows NT in 1992.¹⁰⁴ They renamed the Nexpert Object software "Elements Expert". After C++ came out, they rewrote the product in that language. People built "cool" applications with their product, but nothing was ever deployed, and the business sides of AI firms started to cut off R&D resources.¹⁰⁵

Widespread, But Not A "Killer-App"

Neuron Data's strategy in the late 1980s shifted to designing software tools to be embedded in other applications for a wide variety of platforms, from their original Mac to mainframes. They became the largest seller of AI tools with more than ten thousand licensees.¹⁰⁶ Today, Neuron Data creates "business rules solutions that enable frontline automated systems, such as websites and IVR applications"¹⁰⁷

To have been really successful, commented John Price, the former director of technical marketing at Neuron Data, an AI application needed both to be deployable on PCs, and to be maintainable by a typical IT person. Most early AI firms, including Neuron Data, stumbled in achieving the first, and never came close to the second.¹⁰⁸

4.4 Software Tools

In the software tools category, there were several firms that were building tools for programming in Lisp on multi-purpose computers as opposed to Lisp machines like those from Symbolics or LMI. Notable among those were Gold Hill and Lucid; Gold Hill started with the vision of putting Lisp on PCs, whereas Lucid's offerings put the language on Unix workstations. Lisp never made much of an impact in the corporate world, although its legacy, object oriented techniques, certainly did.

4.4.1 Gold Hill Inc.

"PC Lisp"

In 1984, Gerald Barber (MIT Ph.D.) joined with University of Idaho classmates John Teeter and Stan Curtis to form a company, originally to develop add-in cards for the personal computer, but as that market began to look saturated, they decided instead to bring LOGO and Lisp to the PC. They named the company Gold Hill Computers, after Teeter's property in Idaho. They set up shop on "the wrong side of the tracks" in Cambridge, Mass, in a former Armenian dance hall. They brought on several technical advisors from the MIT AI Laboratory, including Carl Hewitt and Patrick Winston.^{109, 110} In April of that year, they brought out Golden Common Lisp for DOS, the first common Lisp product for the PC. The product also included tutorial software written by Winston.¹¹¹

In 1985, the founders hired Carl Wolf to be President and CEO; Wolf's previous experiences included time at International Data Corporation, strategic planning and venture capital at Xerox; the Boston Consulting Group, and the U. S. Navy. He had become excited about the prospects of PCs and AI, and heard Gold Hill needed a CEO; he saw this situation as an opportunity to get involved in the industry. In an interview, however, Wolf keyed on one of Gold Hill's, and the artificial intelligence industry's, key problems: the market that actually bought their products was very small, because only the very technical could use these products. As Wolf said, "AI products are selling to people who are very, very smart." However, there is a much larger market of people who use technology but are not very, very smart.¹¹²

Products and Funding

By 1986 they had sold more than 6000 copies of GCLisp, mostly for education and low level expert systems development. They also formed an agreement with Symbolics. Although the Lisp computer maker wanted to encourage sales of its own machine for Lisp development use, it realized that the Gold Hill product could give some users a way to "try out" a less powerful version of the technology before investing larger amounts of money into a full-fledged Symbolics system. It was also a way for Symbolics customers who needed multi-terminal systems but were not willing to purchase multiple Symbolics machines a cheaper way to extend their systems, by having a single Symbolics machine acting as a server connected to multiple PCs running GCLisp.¹¹³

Over this period, Gold Hill received funding from BMW Technology Partners Fund, Charles River Partners, Fairfield Venture Management Co., Memorial Drive Trust, and Churchill International.¹¹⁴ However, the Lisp market began to disintegrate, and Gold Hill's fortunes went down with it. It came close to merging with ExperTech, an English AI company, but the deal fell through. In March of 1990, due to pressure from their venture capitalists, they had to lay off more than half its employees (eleven of the twenty-one who had been employed, down from eighty-five at its height).¹¹⁵ The founders, Barber, Teeter and Curtis, each left to look for greener pastures.¹¹⁶

Resuscitation

A few employees came to its rescue: they bought the company, renamed it Gold Hill (from Gold Hill Computers), and former COO Celia Wolf (no relation to Carl) took over as CEO. They continued to develop new versions of GCLisp and GoldWorks, and the company started posting a profit.¹¹⁷ Today, Gold Hill produces "dynamic, object oriented software tools and applications" that include Lisp programming environments and tutorials as well as an expert system development environment.¹¹⁸

4.4.2 Lucid Inc.

"Perpetual Promise" in the AI Languages Industry¹¹⁹

In 1983, John Kulp, then VP of Research and Development at Symbolics, laid off a group of programmers in a cost-saving move. These programmers were working on a Lisp product for mainstream hardware systems. These programmers joined with Tony Slocum from IntelliCorp and formed a new company called Lucid, with Slocum as CEO¹²⁰, in Menlo Park, CA.¹²¹ They joined with Stanford Lisp guru Richard Gabriel to make a product that would enable Lisp to run on Sun and Apollo workstations, so that these "generic" workstations could run like Lisp machines but cost much less than the specialized Symbolics or LMI workstations.¹²² Their plan required Lisp to take off among computer programmers from its more specialized position, however, and that never happened.¹²³

From the start, Lucid was a technology driven company, not a market driven one. Their customers were mainly research laboratories and DARPA projects.¹²⁴ But Lucid was also a venture-backed firm, taking \$15 million over its lifetime from eight venture capital firms,¹²⁵ and the venture capitalists wanted to see venture returns. Since the VCs expect that around 30% of their firms will fail, 50% will fall into the

"walking dead" category, and the rest will be winners, they want to focus their time on the winners.

However, Lucid ultimately ended up as one of the walking dead.¹²⁶

Lucid had a very attractive beginning in putting Lisp on general purpose computers, such as Sun, SGI, and HP workstations. It was structured to be an OEM, providing software for the hardware firms to bundle with their wares. Thus the firm did not need to focus on marketing or sales except to HP, and the other hardware firms. The hardware makers' sales forces sold the Lucid software. This setup was very attractive to Lucid, as a technology focused firm; ultimately IBM, HP, SGI, Sun (most of "the majors" in the hardware industry) all had made OEM agreements with Lucid.¹²⁷

A New Sales Approach

There was a belief among Lucid and its OEM partners that the product would sell in great numbers, but this growth never materialized; \$5 million was Lucid's high point in revenue. Technical issues were never the problem; instead, the majors realized the Lisp market was not growing enough, and they began to shift away from OEM deals and thus Lucid had to switch to a direct sales business model. The Lisp group had to become more involved in the selling and maintenance of its products, but this business model was less attractive to the firm. Lucid concluded that this plan was not going to be very successful, and they decided to find another product that would leverage their knowledge that could provide a better return.¹²⁸

This new product, introduced in 1992, was "Energize": a programming environment for C/C++ languages that offered the integrated environment and tools for object oriented programming that normally were only found in specialized languages like Lisp.¹²⁹ While technically an interesting product, the Lisp user community was actually very different from the C/C++ user communities. The product was not well received in the C/C++ world.¹³⁰ Although it sold reasonably well, it was difficult to install and took a long time to learn. "The mistake we made was trying to solve the world's problems with one product," stated Rigdon Currie, one of Lucid's venture capitalists. The firm ran into a cash crunch when a Japanese distributor, INS, held back on a payment after the layoff of what they felt were key employees, and Silicon Valley Bank became "nervous" about a \$450,000 overdue line of credit. The venture backers were not willing to put up any more money.¹³¹ They thought about splitting the company between Lisp and Energize, which might be able to find VC money to fund a change its marketing strategy. However, making the split between the two parts of the firm proved to be too difficult.¹³²

Selling Off the Pieces

Lucid needed serious cuts to stay alive, and it decided to declare bankruptcy. It told all its creditors that it planned to file Chapter 11, although it never actually did. Its creditors, in the end, got about twenty cents on the dollar, except for the employees, who were paid from the venture capital funds generated during the liquidation of the company. Harlequin Ltd., a UK software company founded in 1986 to solve "hard problems"¹³³, bought the Lisp technologies for sale in everywhere but Japan in what was basically a "fire sale" of Lucid's remaining technologies, hired three to four key people, and still supports the language. Lisp, commented Harlequin's founder and President Jo Marks, still "succeeds where other systems do not."¹³⁴ As for the rest of Lucid, CSK bought the rights to sell the Lisp products in Japan; INS, a subsidiary of Nippon Telegraph and Telephone, bought the rights to sell Energize in Japan; and Tartan bought the rights to license Energize everywhere else.¹³⁵ A competitor, Franz Inc., a small, privately-held object-oriented tools firm founded in Berkeley, CA, in 1984,¹³⁶ was also interested, but the two firms were not able to put a deal together.^{137,138}

In the end, the market failed to materialize for the Lisp language, and Lucid could not meet the expectations of its VC investors. If it had stayed as a closely held technical firm, like Harlequin or Franz, it might still be around, with \$1.5 million in annual revenue,¹³⁹ but it took the chance that the market might grow dramatically. Unfortunately, it did not.

4.5 Expert Systems Applications

Several companies attempted to go after a vertical market by building specific expert system applications. This effort was particularly attractive in the financial services industry. Two firms, Applied Expert Systems and Palladian, saw a potential market in various aspects of finance; for Apex, it was personal financial planning, and for Palladian, operations and financial decision making for corporations. Both companies wrote their products in Lisp and ran them on Lisp-specific hardware, which they soon found was difficult for prospective customers to swallow.

4.5.1 Applied Expert Systems Inc.

Thinking Vertical

Watching the difficulties of Teknowledge and the other expert system shell companies sell their production-rules focused products without facing a barrage of consulting work made at least a few entrepreneurs recognize the potential for building specific vertical expert system applications. In this spirit, Fred Luconi (former MIT computer science professor and founder of Index Systems) and Randy Davis (a founder of Teknowledge and an MIT professor) joined with several others to form Applied Expert Systems (known as Apex) in 1983. They recruited another MIT alumni, Ken Morse, to be in charge of sales and marketing; he stayed for a year in that position and then spent a year as a consultant before moving on. The founders planned to build a specific application, in the financial sector, to take advantage of expert systems technology.¹⁴⁰ However, while the founders realized how to avoid the consulting pitfall of AI, they still fell into another one, that of the specialized platform.

Personal financial planning seemed to be a good business to which to apply this technology. At the time, planners would give their clients fifty page questionnaires in which the client would prepare a list of everything they owned, their expenses, their financial goals and aspirations, and their risk tolerances. Planners would organize this financial information and advise the client on where to invest his money. However, the nation's money was changing hands into the baby boomer generation, and Apex saw a need for more and improved personal financial management services. As the process for advising clients looked like it would fit well into an expert system, they decided to build a product for this purpose.¹⁴¹

Technology Strategy

Luconi described the process of building their products as having two layers of development. First, the company must build a technology robust enough to use in a commercial context, and second, they must build an application on top of it (this second step being the one that most of the other shell firms seemed to be missing). He also envisioned two channels of access. One would be extremely user accessible, in the form of a kiosk in the user's workplace, that would be self-sustaining in terms of usability and software. The second would be for upscale users, such as entrepreneurs, who had very sophisticated and complicated portfolios, with investment, cash flow, etc. all interacting. "Advisory centers" would then, with the help of AI techniques, package their real estate holdings into special products for these customers. This second channel turned into PlanPower; it would pull together taxes, investment, and expertise to leverage special tax breaks and other financial tools for the twenty thousand financial planners in the U.S. Their core

strategy was to create value with AI by solving specialized financial industry problems; and then they would move to a more common platform.¹⁴²

Before jumping into operation, the founders hired an MBA student from Wharton, Arthur Schwartz, to do a feasibility study on its proposed product. The response was mixed: larger firms liked the idea, but smaller firms were worried about the necessary financial outlays. Schwartz went on to join Teknowledge, and Apex decided to go ahead with their product, which they called PlanPower.¹⁴³ Davis took on the role of "Technical Advisor," in which he would train people, help recruit technical hires, and assist in client contact.¹⁴⁴

Teaming with Travelers Insurance

Travelers Insurance turned out to be especially interested in this idea, and they invested an undisclosed sum of money into Apex for them to create this product; when it was finished, Travelers would keep the application, but Apex could use the shell technology base to create other applications. After one and a half years of development, it was clear that PlanPower would work. Part of Apex's organization was dedicated to Travelers, much like a franchise supplier; Apex was responsible for putting PlanPower into different Travelers offices. Apex was able to deliver on this product; since they had not taken any venture capital money, they were able to work under less pressure.¹⁴⁵ They also had two other lesser corporate sponsors.

PlanPower was written in Lisp and built to run on Xerox 1186s, the only machine capable of running their software at that time.¹⁴⁶ In 1985, Xerox signed a deal with Apex to sell 1000 1186s to Apex for \$20 million. Apex would install PlanPower on these machines and sell them to financial planners, particularly those at Travelers, their major client and investor.¹⁴⁷

Market Troubles

However, despite the functionality of the product, Apex was not able to market its downscale version of the product very successfully. Prospective clients balked at having to install specialized hardware, and in this era, most financial planners did not even use computers. The financial planners also had some fear that this software would take their jobs away from them, and this phenomenon led to some distrust of the product's decisions. There was really no public market for their software.¹⁴⁸

They did sell a downscale, broad market version of PlanPower to John Hancock as part of Hancock's "real life, real answers" ad campaign about planning for financial future; Apex delivered a very different

version of the technology to them compared to Travelers. But although they began to move this to an IBM mainframe (a more generic architecture), they never finished doing so. It was becoming clear that software written in specialized languages (like Lisp) for specialized architectures (like the Xerox 1186) were not going to be the wave of the future; that belonged to the PC. Luconi left the company before they ported the software to the PC platform.¹⁴⁹

Unfortunately, with the changes in the tax laws in the 1980s, tax shelters were destroyed by new tax codes. Travelers pulled out of the business of creating shelters and Apex lost them as a client; Apex's revenues dropped from \$8-9 million to \$2 million in a day. Even more so than Carnegie Group, Apex suffered from having a large percentage of its revenues coming from a major client. Luconi saw the company's only future as a consulting business, and in 1989 he left the company to consult for, and then join, Index Technology Corporation.¹⁵⁰ Today, Applied Expert Systems provides network management software based on expert systems technology.¹⁵¹

4.5.2 Palladian Software Inc.

After Phil Cooper sold his first startup, Computer Pictures Corporation, to Cullinet Software in 1982, he decided that his next business should be in artificial intelligence. Enrolling himself as a self-funded Sloan Fellow at MIT's Sloan School of Management, he began learning about applying AI technologies to business problems. His Masters thesis covered a wide range of AI lore, from its history and philosophy to its business applications.¹⁵² In 1984, Cooper was able to use this thesis as a business plan to raise money for his new venture, Palladian Software Inc.¹⁵³

Finding the Experts and Building the Products

Cooper lined up an impressive list of Sloan professors to advise his venture. Furthermore he raised money from some of the top venture capitalists in the industry, including Kleiner Perkins, Venrock, Lazard Freres, Morgan Stanley, and even Harvard University. The firm set up shop in Cambridge above Legal Seafood's restaurant in Kendall Square. Learning from the last set of AI firms, Cooper chose a specific application for AI technologies: he programmed the expertise of the Sloan professors into a product for CFOs called the Financial Advisor for evaluating capital planning opportunities.¹⁵⁴ This product was released in August 1985 at the price of \$95,000.¹⁵⁵

Cooper capitalized on the hype that surrounded AI to publicize his own company. One marketing campaign, in 1987, included buying a billboard in Kendall Square with a sign saying, "The Garden of Eden, Plymouth Rock, Kitty Hawk, Route 128, AI Alley. Great Beginnings.", with the Palladian logo and a picture of himself and his employees.¹⁵⁶ However, his company's fortunes were not that rosy, thanks in part to problems with the product's name, which suggested a much narrower CFO-type market than the general business market the company was aiming for. They updated the product and renamed it the Management Advisor in December 1986. In March 1987, they announced that they ported the financially oriented expert system to the general-purpose Apollo workstation, with the goal of bringing down the price and attracting more customers than the Symbolics machine.¹⁵⁷ The company's revenues were estimated at about \$3 million; presumably, this number was the highest they ever reached.¹⁵⁸

Palladian released its next product, The Operations Advisor, in 1987, for manufacturing strategy planning; the firm charged \$136K for a copy. Palladian had difficulties making headway with the product, however, and sold it to Carnegie Group later that same year for \$30K.¹⁵⁹ Carnegie renamed the product the Operations Planner and by 1990, it ran on a PC and sold for \$4,000.¹⁶⁰

Falling Apart

In August of 1987, after more than \$20 million in venture funding and little in return, Cooper resigned. The company's revenues were estimated at just \$1-2 million per year, and after losing its top three programmers, the firm was having difficulties developing any new products.¹⁶¹ Cooper moved to the other side of the table, working first for Harvard's venture capital group, and he is currently a partner at Goldman Sachs, working in private equity. James McGowan from IBM took over the CEO spot,¹⁶² and after laying off the rest of its workforce and shutting offices, Palladian finally shut its doors in May 1989.¹⁶³

In retrospect, Cooper says that the company was too tied to Lisp; they were not able to get off the Lisp platform in time. Their product needed to be built on a standardized workstation in a common language; the Lisp/Symbolics combination, which was the best available at the time, were too inflexible for the long-term. While competition was not a problem, the price was. Unfortunately it was not obvious at the time that Lisp would become irrelevant in the corporate world; today, its strengths, such as its object-oriented features, have been incorporated into today's languages.¹⁶⁴

4.6 Knowledge Base Firms

One other interesting area within AI software is the common sense knowledge base. In the laboratories, AI researchers debate about the usefulness and the viability of creating a common sense knowledge base. While many agree it would be nice to have, as a basis for building better "expert" systems, most believe it would take so long to encode all the knowledge (provided one could even find the right way to represent it) it would not be a worthwhile commercial venture. Nevertheless, one researcher, Doug Lenat, decided this effort was too important to pass up. For ten years at the Microelectronics and Computer Technology Corporation (MCC), Lenat worked on a common sense knowledge base. Five years ago he spun this project out into a private company, Cycorp. Its story is still far from complete.

4.6.1 Cycorp

Building A Common Sense Knowledge Base

An interesting project to look at in the history of AI companies is that of Doug Lenat's project Cyc (pronounced "psych", from Common Sense Knowledge-Base Construction), which later spun out of the MCC to become the company Cycorp. Lenat had received his Ph.D. in computer science from Stanford University and taught at CMU and Stanford. Cyc began in 1984 at the Microelectronics and Computer Technology Corporation (MCC) research consortium in Austin, TX. The MCC began in response to the Japanese Fifth Generation Project, and was funded by the Pentagon as well as member corporations to do basic high-tech research.¹⁶⁵

Lenat and partner Ramanathan Guha strove to build a common-sense knowledge base using 'ontological engineers' to encode common-sense facts that ideally would one day mimic human common sense. However, projections for the end of this project were continually optimistic; the original estimate was ten years; and in 1989 they guessed it would only take five more years to prime Cyc.¹⁶⁶ Today estimates are around another ten years before they have a consumer product.¹⁶⁷

However, by 1993 things were not looking good for Cyc at the MCC; several of the funders of the project were complaining about the overhead (300%). After a long battle, the MCC decided to terminate Cyc at the end of 1994. Lenat then founded Cycorp in 1995 and licensed the Cyc technology from the

MCC. Guha left the project and is currently at Netscape Communications. Since then, Cycorp has received funding from Apple (in 1995), the NSA, Glaxo, and other government agencies.¹⁶⁸

Today Cycorp continues to develop the Cyc knowledge base and is working to develop applications to take advantage of this knowledge in individual projects. Cycorp is still a technology-focused company, however, and from some reports still struggling to find its market; according to one source, Lenat's hold on the company makes it difficult to recruit a CEO or any marketing or sales people. Coming straight from the research environment, Cycorp does not have the experience of delivering consumer software products and may find it difficult to do so, even if it can get its knowledge base complete.

4.7 AI Software Industry Summary

By the middle of the 1980s, hype about the AI software industry was running high. "The cash is flowing," commented one 1985 Business Week article.¹⁶⁹ While much of Silicon Valley was slumping during the same year, artificial intelligence was one of the few bright spots; AI and genetic engineering, both of which peaked and plummeted.¹⁷⁰

The hype for expert systems among the corporate world was impressive. A 1986 study commissioned by Coopers & Lybrand indicated that although most financial services executives were "less than enthusiastic" about the future of expert systems, a "significant percentage" said their strategic potential was such that they could not afford not to be involved.¹⁷¹

But such hype spelled trouble for the industry. Expert systems were difficult to build; determining which functions would most benefit from such a system, encoding that information into a system, and tying that system into internal databases were all difficult tasks. Although there were several well-publicized success stories (XCON, American Express), many more systems were put together in prototype but never made it to the production stage.

At the end of the 1980s, when the reality caught up with the hype, the firms' underlying problems (management, marketing, technology) came to bear heavily upon them. By then the press had started to declare the "AI Winter," and for many of the firms, it was too late to change their ways without undergoing massive restructuring. However, the AI industry did realize that it was not, in fact, an industry; today there are customer relationship management firms, and enterprise resource planning firms, using AI

technologies, but generally there is little mention of an AI industry. The image of the technology has matured and become another tool in the programmer's toolbox.

-
- ¹ John Price, Interview by the author, Austin, TX, 9 March 1999.
 - ² Matt Rothman and Emily Smith, "The Leading Edge of 'White-Collar Robotics' – These Hot Startups Are Rushing to Cash in on Computer Software That Mimics Human Reasoning," *Business Week* (10 February 1986), 94.
 - ³ Richard Duda and Edward Shortliffe, "Expert Systems Research," *Science* (15 April 1983), 261.
 - ⁴ Duda and Shortliffe, *Science*.
 - ⁵ Duda and Shortliffe, *Science*.
 - ⁶ Duda and Shortliffe, *Science*.
 - ⁷ Microsoft Encarta Encyclopedia, Web Edition, "Gang of Four."
 - ⁸ Rothman and Smith, *Business Week*.
 - ⁹ Joyce Gannon, "Perq Lays Off 30; Spins Off AI Firm," *Pittsburgh Business Times & Journal* (28 January 1995), 1.
 - ¹⁰ HP Newquist, *The Brain Makers: Genius, Ego and Greed in the Quest for Machines that Think* (Indianapolis, Sams Publishing, 1994), 304-305.
 - ¹¹ Newquist, *The Brain Makers*, 307.
 - ¹² Brian Cook, "AI: Industry's New Brain Child" *Industry Week* (1 April 1991), 54.
 - ¹³ "Carnegie Group Inc. Integrates Its Two Software Products, 'Knowledge Craft' and 'Language Craft'", *PR Newswire* (28 April 1986), 402.
 - ¹⁴ Newquist, *The Brain Makers*, 356.
 - ¹⁵ Seth H. Lubove, "Carnegie Group President's Resignation Shows Tension at Some Start-Up Firms", *The Wall Street Journal* (20 February 1987).
 - ¹⁶ Joyce Gannon, "Start-Up Slowed: Software Firm Says Loan Dispute Caused Bankruptcy Filing," *Pittsburgh Post-Gazette* (19 September 1991), 14.
 - ¹⁷ Paul Sagawa, "The Carnegie Group" (Harvard Business School case study; originally published 1990; revised 11/10/92; #9-690-033).
 - ¹⁸ Newquist, *The Brain Makers*, 355.
 - ¹⁹ Peter Burrows, "Artificial Intelligence Finds Work in Manufacturing," *Electronic Business Buyer* (5 February 1990), 65.
 - ²⁰ Sagawa, HBS case study.
 - ²¹ Kendra Berteotti, "High Hopes at Carnegie Group for New Initiative," *Pittsburgh Business Times & Journal* (30 April 1990) 2.
 - ²² Barbara Klein, "A Risk-Taker Sizes Up the Rewards," *Pennsylvania Business and Technology* (1 April 1992) 32.
 - ²³ "Carnegie Group Signs Contract with Caterpillar for Development of a Fully Automatic Text Translation System," *PR Newswire* (8 May 1992).
 - ²⁴ "Carnegie Group Announced the Availability of Rock 2.1," *PR Newswire* (7 July 1992).
 - ²⁵ "Carnegie Stock Up," *Pittsburgh Post-Gazette* (5 December 1995), c-9.
 - ²⁶ Karen Kovatch, "Carnegie Group Seeks New Clients to Bolster Investors," *Pittsburgh Business Times & Journal* (16 June 1997), 5.
 - ²⁷ Karen Kovatch, "Unfulfilled Promise?" *Pittsburgh Business Times & Journal* (23 October 1998), 1.
 - ²⁸ "Logica to Acquire Carnegie Group for \$35 Million," Logica plc Press Release (1 October 1998).
 - ²⁹ Newquist, *The Brain Makers*, 186-187.
 - ³⁰ Newquist, *The Brain Makers*, 263-264.
 - ³¹ "IntelliGenetics' 1.6 Million Initial Shares Offered at \$6," *Dow Jones News Service - Ticker* (15 December 1983).
 - ³² Newquist, *The Brain Makers*, 292.
 - ³³ Ken Ross, Interview by the author, Mountain View, CA, 13 January 1999.
 - ³⁴ "IntelliCorp Sells Remaining Interest in Joint Venture to Amoco," *Business Wire* (9 November 1990).
 - ³⁵ "Oxford Molecular Acquires IntelliGenetics, Inc." *PR Newswire* (25 August 1994).
 - ³⁶ Bob Fondiller, Interview by the author, Sunnyvale, CA, 9 February 1999.
 - ³⁷ Ross, 13 January 1999.

-
- ³⁸ "Sperry Corp. Signs Contract to Purchase Texas Instruments' Explorer Systems as Part of Its Knowledge System Product Line," *PR Newswire* (25 June 1985), 66.
- ³⁹ Newquist, *The Brain Makers*, 308-311.
- ⁴⁰ Douglas Keith, Interview by the author, Boston, MA, 21 January 1999.
- ⁴¹ Jim Schembari, "Firm Positioned to Gain From New Technology," *Chicago Sun-Times* (17 November 1985), 4.
- ⁴² Gregory Spector, "Interface Helps Build Knowledge-Based Systems," *PC Week* (29 April 1986 v3), 18.
- ⁴³ William Bulkeley, "Stocks of Artificial Intelligence Firms Prosper, Though Some Analysts Are Advising Wariness," *The Wall Street Journal* (31 March 1986).
- ⁴⁴ Jeff Moad, "Building a Bridge to Expert Systems; IntelliCorp Claims It's Found a Way to Link Its Knowledge Systems with Relational Databases," *Datamation* (1 January 1987 v33), 17.
- ⁴⁵ "Expert-system Development From IntelliCorp Now Runs on PCs," *PC Week* (5 April 1988, v5 n14), 26.
- ⁴⁶ Newquist, *The Brain Makers*, 421-423, 435.
- ⁴⁷ IntelliCorp SEC Annual Report (Form 10-KSB), For Fiscal Year Ended June 30, 1998.
- ⁴⁸ Newquist, *The Brain Makers*, 225.
- ⁴⁹ Chuck Williams, Interview by the author, Novato, CA, 23 February 1999.
- ⁵⁰ Williams, 23 February 1999.
- ⁵¹ Newquist, *The Brain Makers*, 401.
- ⁵² Williams, 23 February 1999.
- ⁵³ "Lockheed to Assume Minority Position in Inference Corp," *Dow Jones News Service - Ticker* (21 December 1984).
- ⁵⁴ "Lockheed to Invest Additional \$2 Million in Inference Corp.," *Dow Jones News Service - Ticker* (4 March 1986).
- ⁵⁵ "Ford Acquires 10% Stake in L.A. Software Company," *Los Angeles Times* (25 October 1985), 2.
- ⁵⁶ Williams, 23 February 1999.
- ⁵⁷ "Inference Closes \$5.5 Million in Growth Financing; New Round of Funding Led by JP Morgan," *Business Wire* (15 July 1991).
- ⁵⁸ Williams, 23 February 1999.
- ⁵⁹ Newquist, *The Brain Makers*, 275-276.
- ⁶⁰ Newquist, *The Brain Makers*, 401-402.
- ⁶¹ Esther Dyson, "Your Humble Servant (Artificial Intelligence Systems)," *Forbes* (2 November 1987), 204.
- ⁶² Williams, 23 February 1999.
- ⁶³ Williams, 23 February 1999.
- Darryl K Taft, ⁶⁴ "Art-ADA Marries Expert Systems, Ada Technology," *Government Computer News* (8 January 1990), 33.
- ⁶⁵ Paula Musich, "Software Upgrade, AI Program Boost AS/400's PC Connectivity," *PC Week* (11 September 1989), 6.
- ⁶⁶ Williams, 23 February 1999.
- ⁶⁷ Williams, 23 February 1999.
- ⁶⁸ Williams, 23 February 1999.
- ⁶⁹ Williams, 23 February 1999.
- ⁷⁰ Newquist, *The Brain Makers*, 219-221.
- ⁷¹ Newquist, *The Brain Makers*, 265-266.
- ⁷² Newquist, *The Brain Makers*, 295-298.
- ⁷³ Randy Davis, Interview by the author, Cambridge, MA, 20 October 1998.
- ⁷⁴ John Halusha, "General Motors Corp, in Little-Noticed Move, Pays \$3 Million for 11 Percent Interest in Teknowledge Inc.," *New York Times Abstracts* (3 May 1984), 1.
- ⁷⁵ Newquist, *The Brain Makers*, 307.
- ⁷⁶ Udayan Gupta, "Concerns Rush To Go Public As Technology Issues Prosper," *The Wall Street Journal* (24 March 1986).
- ⁷⁷ Robert Sonenclar, "Going Public; The New Issues of 1986," *FW* (18 March 1986), 20.
- ⁷⁸ Newquist, *The Brain Makers*, 298.
- ⁷⁹ Newquist, *The Brain Makers*, 334.

-
- ⁸⁰ William Bulkeley, "Stocks of Artificial Intelligence Firms Prosper, Though Some Analysts Are Advising Wariness," *The Wall Street Journal* (31 March 1986).
- ⁸¹ Newquist, *The Brain Makers*, 365-366.
- ⁸² Newquist, *The Brain Makers*, 367, 380-381, 387-9.
- ⁸³ Teknowledge Website, "Corporate Information" (<http://radon.teknowledge.com/overview/index.html>).
- ⁸⁴ William Pat Patterson, "Natural language lets anyone use computers," *Industry Week* (2 April 1984 v221), 62.
- ⁸⁵ Patterson, *Industry Week*.
- ⁸⁶ Jan Johnson, "Easy Does It: Natural Language Query Systems Have Gained New Attention in Recent Months," *Datamation* (15 June 1984 v30), 48.
- ⁸⁷ Paul Gillin, "Plans Dropped for PC Version of 'Intellect,'" *PC Week* (5 March 1985 v2), 123.
- ⁸⁸ Barbara Buell, "The Professor Getting Straight As on Route 128 - Larry Harris Mothballed His Cap and Gown to Start a Company that Cashes In on Artificial Intelligence," *Business Week* (15 April 1985), 150J.
- ⁸⁹ Ralph Emmett Carlyle, "Sneaking In the Back; Get Ready For a Quiet Influx of Expert Systems Products From Several Familiar Faces," *Datamation* (1 October 1986 v32), 32.
- ⁹⁰ John Desmond, "When MIS and AI Converge," *Software Magazine* (1 May 1988), 58.
- ⁹¹ John Parker, "Merger Reflects New Focus for Expert Systems," *PC Week* (2 April 1990 v7 n13), 116.
- ⁹² Lawrence Edelman, "AICorp of Waltham Goes Public Sale of 2.25 Million Common Shares \$21.4 Million," *The Boston Globe* (27 June 1990), 45.
- ⁹³ "AICorp and Aion Sign Merger Agreement," *Business Wire* (25 June 1992).
- ⁹⁴ "AICorp and Aion Sign Merger Agreement," *Business Wire* (25 June 1992).
- ⁹⁵ "PLATINUM technology and Trinzic Complete Merger," *Platinum News Release* (25 August 1995).
- ⁹⁶ PCAI Interview: Patrick Perez, *PC AI Magazine* (Nov/Dec 1989), 34.
- ⁹⁷ PCAI Interview: Patrick Perez, *PC AI Magazine* (Nov/Dec 1989), 34.
- ⁹⁸ John Price, Interview by the author, Austin, TX, 9 March 1999.
- ⁹⁹ Newquist, *The Brain Makers*, 365.
- ¹⁰⁰ Newquist, *The Brain Makers*, 420.
- ¹⁰¹ Pat Bellamah, "Neuron Enhances Its Expert System to Link to Many Databases," *PC Week* (25 August 1988 v5 n30), 33.
- ¹⁰² John Pallatto, "Expert-system Shells Span DOS, OS/2 Nets," *PC Week* (28 August 1989 v6 n34), 6.
- ¹⁰³ Marc Ferranti, "Neuron readies more flexible Nexpert," *PC Week* (2 July 1990 v7 n26), 43.
- ¹⁰⁴ "Neuron Data Inc (announced support for Microsoft Corp.'s Windows NT operating system)," *InfoWorld* (23 November 1992 v24 n47), 15.
- ¹⁰⁵ Price, 9 March 1999.
- ¹⁰⁶ Newquist, *The Brain Makers*, 442.
- ¹⁰⁷ Neuron Data Website, "Company Information" (<http://www.neurondata.com/Company/index.html>).
- ¹⁰⁸ Price, 9 March 1999.
- ¹⁰⁹ Paul Gillin, "Lisp Packages Bespeak Success for AI Tools Firm," *PC Week* (1 April 1986 v3), 101.
- ¹¹⁰ Newquist, *The Brain Makers*, 279.
- ¹¹¹ Richard Duffy, "Lisp Language for the PC Introduced; Gold Hill Says GCLisp Supports DOS, Cost \$495," *PC Week* (17 April 1984 v1), 4.
- ¹¹² "An Interview with Carl Wolf," *PC AI Magazine*, (Sept/Oct 1988).
- ¹¹³ Paul Gillin, "Lisp Packages Bespeak Success for AI Tools Firm," *PC Week* (1 April 1986 v3), 101.
- ¹¹⁴ R. Colin Johnson, "Layoffs Hit AI Specialist Gold Hill," *Electronic Engineering Times, CMP Electronics File* (12 March 1990), 15.
- ¹¹⁵ Johnson, *Electronic Engineering Times, CMP Electronics File*.
- ¹¹⁶ Newquist, *The Brain Makers*, 418-419.
- ¹¹⁷ Ellen T. Corliss, "Gold Hill Fights the Odds," *Mass High Tech* (30 July 1990), 1.
- ¹¹⁸ Gold Hill Website, Background Information, <http://www.goldhill-inc.com>
- ¹¹⁹ "A Call for Candor," *Venture Capital Journal* (1 March 1995).
- ¹²⁰ Newquist, *The Brain Makers*, 262, 292.
- ¹²¹ Renee Deger, "Blinded By Science," *Venture Capital Journal* (1 March 1995).
- ¹²² Newquist, *The Brain Makers*, 331.
- ¹²³ Deger, *Venture Capital Journal*.
- ¹²⁴ Newquist, *The Brain Makers*, 292.

-
- ¹²⁵ Deger, *Venture Capital Journal*.
- ¹²⁶ John Tannenbring, Interview by the author, Saratoga, CA, 8 February 1999.
- ¹²⁷ Tannenbring, 8 February 1999.
- ¹²⁸ Tannenbring, 8 February 1999.
- ¹²⁹ "Lucid launches Advanced C++ and C Programming System for SPARC; Provides a Natural Path to Object-Oriented Programming," *Business Wire* (2 January 1992).
- ¹³⁰ Tannenbring, 8 February 1999.
- ¹³¹ Deger, *Venture Capital Journal*.
- ¹³² Tannenbring, 8 February 1999.
- ¹³³ Harlequin Website, "About Harlequin" (<http://www.harlequin.com/about>).
- ¹³⁴ "Interview: Jo Marks," *PC AI Magazine* (Jan/Feb 1995), 39.
- Renee Deger, "Blinded By Science," *Venture Capital Journal* (1 March 1995).
- ¹³⁶ Franz Website, "About Franz" (<http://www.franz.com>).
- ¹³⁷ Tannenbring, 8 February 1999.
- ¹³⁸ Simon Clephan, Email received by the author, 30 January 1999.
- ¹³⁹ Tannenbring, 8 February 1999.
- ¹⁴⁰ Fred Luconi, Interview by the author, Florida, 10 February 1999.
- ¹⁴¹ Randy Davis, Interview by the author, Cambridge, MA, 20 October 1998.
- ¹⁴² Luconi, 10 February 1999.
- ¹⁴³ Newquist, *The Brain Makers*, 295.
- ¹⁴⁴ Davis, 20 October 1998.
- ¹⁴⁵ Luconi, 10 February 1999.
- ¹⁴⁶ Davis, 20 October 1998.
- ¹⁴⁷ Newquist, *The Brain Makers*, 334.
- ¹⁴⁸ Davis, 20 October 1998.
- ¹⁴⁹ Luconi, 10 February 1999.
- ¹⁵⁰ Luconi, 10 February 1999.
- ¹⁵¹ "Systems Integration Report: Expert System Cleverly Diagnoses SNA Connectivity Problems," *Manufacturing Automation* (1 December 1994).
- ¹⁵² Cooper's Masters Thesis.
- ¹⁵³ Newquist, *The Brain Makers*, 313-315.
- ¹⁵⁴ Newquist, *The Brain Makers*, 314-320.
- ¹⁵⁵ William M. Bulkeley, "Computers Take On New Role As Experts in Financial Affairs," *The Wall Street Journal* (7 February 1986).
- ¹⁵⁶ Newquist, *The Brain Makers*, 350-355.
- ¹⁵⁷ Michael R. Ball, "Palladian & Apollo Launch 'Bargain' AI System," *Boston Business Journal* (30 March 1987), 14.
- ¹⁵⁸ Ronald Rosenberg, "Palladian Software's President Resigns," *The Boston Globe* (27 August 1987), 57.
- ¹⁵⁹ Newquist, *The Brain Makers*, 355.
- ¹⁶⁰ Peter Burrows, "Artificial Intelligence Finds Work in Manufacturing," *Electronic Business Buyer* (5 February 1990), 65.
- ¹⁶¹ Michael Ball, "Ring-Around-the-Rosy: What's Behind the Recent Shake-Up?" *Boston Business Journal* (7 September 1987), 1.
- ¹⁶² Ball, *Boston Business Journal*.
- ¹⁶³ Newquist, *The Brain Makers*, 355.
- ¹⁶⁴ Phil Cooper, Interview by the author, New York, NY, 19 November 1998.
- ¹⁶⁵ Mark Clifford, "Artificial Intelligence (Investing In High-Tech Firms)," *FW* (23 January 1985), 13.
- ¹⁶⁶ Kevin Kelly, "A High-Tech Think Tank Thinks Big Bucks," *Business Week* (25 September 1989), 222.
- ¹⁶⁷ Ramanathan Guha, Interview with the author, 24 February 1999.
- ¹⁶⁸ Guha, 24 February 1999.
- ¹⁶⁹ Emily Smith, "A High-Tech Market That's Not Feeling The Pinch – Eager Investors Have Created a Boom in Artificial Intelligence, But Can All These Startups Deliver?" *Business Week* (1 July 1985), 78.
- ¹⁷⁰ Daniel Rosenheim, "Silicon Valley Slump – It's Not All That Bad," *The San Francisco Chronicle* (23 August 1985), 6.
- ¹⁷¹ "The Future of Technology in the Financial Services Industry," *American Banker* (14 April 1986), 14.

5 The AI Hardware Industry

In 1974, Richard Greenblatt, Jack Holloway, and Tom Knight were researchers at the MIT Artificial Intelligence Laboratory, which at the time was almost entirely DARPA funded. The machines they worked with there were a far cry from today's norms. The standard computers were "big, time-sharing monsters"; the standard interfacing technology between them was through Teletypes and character-only display screens, which allowed only 24x70 characters, often in uppercase only. There were no mice; clock rates ran about 300-500 KHz, about a factor of 1000 times slower than today's processors. Networking barely existed, and certainly not Ethernet or LANs; the AI lab itself had one of the earliest prototypes. A large disk capacity was 10 megabytes; primary memory had 18 bits of address space. This primary memory was very important for AI applications: 1 megabyte was the amount one could conceivably address, due to the virtual memory limit (memory cost about 1-2 cents per bit). In the words of Tom Knight, things were "unimaginably primitive".¹

Knight and Greenblatt set out to change this state of affairs. Programming languages, such as Lisp, were already reasonably sophisticated, but traditional machines could not run its programs efficiently. Currently the hardware that ran their programs was stifling development. In consequence of certain features of Lisp, Lisp programs tended to be large and thus would occupy a great deal of core memory by the standards of the day. In his thesis, Knight describes his motivations for building the Lisp machine as inadequate virtual address space for large user programs; inadequate computing power for development if using intelligent programming tools, and inefficient information coding of compiled instructions.² They thus decided to build a machine specifically designed to run symbolic programs.³ That was not all, however; they were not just going to change the hardware. While current hardware technology also included excellent timesharing and debugging technologies, and a decent development environment, it also lacked intuitive user interfaces.⁴

By 1974, they had a prototype of this machine; in 1976 they had built the second generation, and produced half a dozen of these machines. They became the workhorse machines for Lisp software development at the AI Laboratory; over time, the group at MIT built another 25 machines. In 1980 it became obvious that other people wanted this machine as well; about fifteen to twenty people from MIT

were already spending most of their time on machine development. They had requests from Xerox Park, CMU, Stanford; but MIT was not in the business of making computers for other labs, and was not about to build them.⁵

Two companies were founded on this potential market, Lisp Machines Inc. (LMI) by Greenblatt, and Symbolics. The rest of this chapter will provide an in-depth look at these two companies as well as activities by TI and Xerox.

5.1 Lisp Machines Inc.

Lisp Machines Inc. (LMI) was founded in 1980 in Cambridge, Massachusetts, by Stephen Wyle, formerly a mid-level manager at his father's Wyle Laboratories,⁶ Richard Greenblatt, and Alex Jacobson, a consultant from Inference Corp. currently working with Control Data Corporation. Jacobson took a seat on the board, and Wyle and Greenblatt ran the company. LMI had a license to use MIT's Lisp processor,⁷ called the CADR after a common Lisp command, and developed by Greenblatt and Tom Knight, to build their own Lisp machines.

The first year, they sold about a dozen machines at \$100,000 each.⁸ But by the end of 1983, they had still only sold a total of twenty-three machines, each of which was built by hand, and it was clear that Symbolics, with its superior cash flow, was moving ahead.⁹ In May of 1983, Frank Spitnogle joined LMI as President and COO after spending 16 years with TI.¹⁰ They realized they needed some outside investment, as distasteful as this was to Greenblatt, and, in 1983, with Wyle as President, Texas Instruments acquired 25% of LMI.¹¹ They also took on further funding from traditional venture capital sources.

Moving Beyond the MIT Design

After shipping twenty-five of the CADR machines, their next product, the Lambda, was ready for market by September of 1983. The Lambda combined the CADR-like processor (based on TI's Nu machine, acquired from Western Digital) with a Motorola Unix processor, the 68010, connected by a high-speed bus. They also were the sole distributor of LM-Prolog, a version of Prolog written in Lisp microcode.¹²

In July 1984, they planned to ship the Lambda 2x2, which would be the first two-user Lisp machine, and the 2x2 Plus, which offered Unix.¹³ This product was followed by the Lambda 4x4 (for four users).


LMI also began extending the software side of its business, seeing the synergy between selling hardware and expert systems. They brought out a process-control expert system, PICON, and announced marketing agreements with process-control firms like Honeywell and Leeds & Northrup in 1985.¹⁴

However, by the end of 1984, TI was tired of dealing with LMI's manufacturing problems and decided to build its own machines. They licensed the Lisp design from LMI (who had little choice but to agree to the deal). TI won a deal to sell 400 Lisp machines to the MIT Laboratory for Computer Science (at extremely low prices that Symbolics and LMI could not match), and their place in the Lisp machine industry was established.¹⁵

Trying to Compete with TI and Symbolics

However, LMI began losing much of its market share to Symbolics in 1985.¹⁶ The firm had reached 170 people, and had not been profitable since 1980. In the summer of 1985, Wyle brought in a former group VP at Digital Equipment Corp., Ward MacKenzie, to be CEO and Chairman of LMI. By November of that year MacKenzie had brought in \$18.1 million more in venture capital funding, which brought their total investment to \$40 million¹⁷ for a company that had tallied 1984 sales of \$10 million.¹⁸ The board decided to move most of the company from its offices in Cambridge and Los Angeles to the less-expensive Andover, Massachusetts.¹⁹

Table 5.1: Estimated LMI Revenues²⁰



LMI began to look more at its software business, and the PICON product it developed, but in the summer of 1986 the PICON group left to form their own company, Gensym, to build process control expert systems. There was little LMI could do, however, because they were in the process of filing for bankruptcy, and the entire firm had moved back to the small office they had kept in Cambridge.²¹ In that same year, the dedicated Lisp machine market was \$200 million in size; it was split by Symbolics with a 58% share; LMI with 13%; TI, 12%; Xerox, 9%; and a smattering of others with the remaining 8%.²²

LMI filed for bankruptcy on April 1, 1987; MacKenzie left to head up Data General's OEM group. In June, Gigamos, the US subsidiary of a Canadian company, purchased its remaining assets. Gigamos then became embroiled in a lawsuit on an unrelated issue from a Japanese firm, and the lawsuit sunk Gigamos. Even the remains of LMI were gone.²³

Gensym, however, made it through the tumultuous AI 1980s and went public in 1996, although its stock today languishes at about a quarter of its original offering price. The company provides software and services for “intelligent operations management. Common applications include quality management, process optimization, dynamic scheduling, network management, energy and environmental management, and process modeling and simulation.”²⁴ It currently has over 250 employees, annual revenues of \$35 million and with a market capitalization of about \$20 million.

5.2 Symbolics, Inc.

Symbolics was officially founded in 1980 by Russell Noftsker and Robert Adams, a lawyer who left the company in 1981. Once they raised \$250,000 from the classic startup fund of friends, family and fools, fourteen AI lab hackers, including Knight, joined them. They later raised \$1.5 million from various prominent East Coast venture capitalists, and in early 1982, they received \$500,000 from General Instruments.²⁵

From the start, there was a business and cultural rift between the two companies. Symbolics took traditional venture capital funding early on, and in multiple rounds. Greenblatt wanted LMI to bootstrap itself and reinvest its profits in development; nevertheless, it eventually took at least as much venture capital money as Symbolics. Bootstrapping, as LMI was trying to do, proved not to be a good route for fast growth or effective product development.²⁶

The Birth of the Lisp Machine Industry

All three firms were building basically the same as Knight's machine; LMI sold multiple dozens. Neither LMI nor TI did much significant architectural work; they did not innovate on the instruction set. Even TI's single chip product used Knight's design.²⁷ This machine was a marvel to behold with 1980 eyes: it had windows, a bitmapped display, and used mice for its user interface. Development tools, including garbage collector, were built into the system; and the documentation was excellent, available both in printed form and in a hypertext online form, much like today's web sites.²⁸

Symbolics was headed by Noftsker, who had been the administrative officer of the AI Laboratory, though he had left MIT by the time of the company's founding. The company was bi-coastal from the start, with the Cambridge, Massachusetts group doing software and architecture development and design, and the Los Angeles group (in Chatsworth) doing production engineering and manufacturing. This system worked out remarkably well, even better than expected; the two groups had drawn clear lines of roles and responsibility. Noftsker had manufacturing contacts on the West Coast and thus decided to build that facility there.²⁹

Preparing for the Next Generation

On the first day of work at Symbolics, the team started on the company's new machine, the 3600. It would move from a word length of 32 bits to 36 bits; the new machine had increased virtual memory as well. Meanwhile, they produced the old machine in order to have a revenue stream; in those days, companies were expected to make money from the start, and be profitable when they went public.³⁰

The 3600 prototype, built in Cambridge, was originally planned to ship in July of 1982. However, the machines were still not available by November. By this time, Noftsker needed more cash from the venture capitalists to keep the company afloat but could not get funding without a working product. He sent teenage programmer Howard Canon to Los Angeles to make 3 3600 machines work, which he did, and received a Porsche for his troubles. Another employee, Robert Strauss, put up enough cash to tide the company over until they received their round of investment, \$2.6 million.³¹

The 3600, once developed, was a very well-engineered system. There were several different incarnations, and at least three fairly major implementations, even one in gate array ASIC form (from the West Coast) However, the Symbolics research team continued to suffer from focus and direction; their efforts diverged into displays, printers, and other peripherals, as they strove for a "clean design" above all

other goals, such as low cost, on-time delivery, or customer needs.³² The East Coast, meanwhile, was working on the integrated Ivory processor. The 3600 was the mainstay of the company, about 1000-2000 machines were built; all of which were made on the West Coast, so that the East Coast could design in peace and not be distracted by the manufacturing side. Symbolics was able to keep a generation or two ahead of their major competitors, TI and LMI.³³

Growing Pains

By the beginning of 1984, Symbolics had reached 300 employees. The venture capitalists thought Noftsker was spending too much, and he left the CEO position in April, though he stayed on as Chairman. Former VP of Research and Development John Kulp was made the COO. Kulp started the selling process for the next machine, the 3670s, but they were not ready to ship yet. Unfortunately this action killed off orders for the 3600, as customers were waiting for the next generation product. The board, worried about revenues, brought Noftsker back in and Kulp left. Noftsker shipped the not-working 3670s to get sales on books, and then had them fixed afterwards.³⁴

In another management mishap, the company brought Bruce Gras into company as the sales and marketing leader. Gras was a former Shearson employee who had brought in Gold Hill to create an expert system called K:Base to do interest rate swaps. However, when Symbolics sent out its IPO prospectus in November of 1984, they found out Gras was in bankruptcy; he was promptly removed. Symbolics still went public that November.³⁵

The hype in the corporate world about AI caused a lot of sales of Symbolics machines to people who did not know what to do with them. However, the company found it difficult to hold back on the fast growth. The strain was beginning to show by the middle of the 1980s; TI, which could afford to lose money on its Lisp machines, was putting price pressures on LMI and Symbolics. The Unix community also started putting price pressure on the company; a Symbolics machine was \$25-30,000, whereas the Sun 6800, which by its specification sheet looked very competitive, was much cheaper and could be paired with some Unix Lisp software. In reality, the systems were hugely different in productivity measures: Symbolics was a much better machine but also a factor of two costlier.³⁶

Hitting the Wall: Management Roulette

By 1987, it was becoming clear that everyone who wanted a Lisp machine had one, and the price of their machine was dropping as well, from \$140,000 to \$40,000 in six years. Noftsker laid off people and moved everyone back to Cambridge.³⁷ Nevertheless, In 1988 Symbolics released the MacIvory, a version of their Ivory chip that offered direct-execution engine for the Macintosh.³⁸

However, these new directions were causing friction within the management team, such that Brian Sear, the company's President, thought they should be focusing on current products, and Noftsker wanted to work on new hardware designs. They took their problems to the board, and they both decided to resign. Jay Wurts became the CEO in 1988 and focused the firm on the Ivory and Genera software tools, and took the company into graphics. He moved the company to Burlington, Massachusetts from Cambridge, and began to streamline and refocus the company's operations.³⁹ He reduced the workforce to 450 people and tried to make good on the over \$70 million in venture funding the firm had received.⁴⁰


Scrambling to Survive

Outside of their technology difficulties, Symbolics was over-extended in the real estate market. In 1986, they had signed ten-year leases on office space in Burlington, Massachusetts, and Los Angeles in anticipation of further growth, but it turned out they now did not need it and they were having trouble subletting it. After some restructuring, however, they continued to bring out new products. In 1989, Symbolics announced the Ivory Lisp co-processor board for various Sun Unix workstations, which integrated Symbolic's Genera, a proprietary AI operating system, with the SunOS Unix-based operating system.⁴¹ By this time their software armory included, as well as Genera, Plexi, a development environment for designing neural networks; Joshua, an ES programming language and environment, Statice, an OO database, Concordia, a hypertext document management system, and Close, an execution environment for the Intel 386-running Unix. The company began to show profits again.⁴²

These efforts were not enough to save the firm, and in fact, their scattered nature may have hurt the firm just when it needed to focus. In 1991, the board brought in Kenneth Tarpey as President and CEO. Tarpey sold off various pieces of the firm and focused the rest on software tools. They even started re-writing their Lisp products in C++.⁴³ Bad real-estate deals continued to haunt them, however, and Symbolics filed for bankruptcy on January 28, 1993.⁴⁴

From the start, commented Knight, who is still at the MIT AI Laboratory, Symbolics desperately needed better management; they never had better than mediocre management and marketing. Even with better marketing and management, Knight said, the company probably would have struggled. Facing pressure from the Sun workstation market and the decline in government funding and purchasing, Symbolics had no place to go. However, they had good programming languages, development environments, and intelligent programs: they were just good ideas waiting to happen.⁴⁵ And today, the Symbolics environment can run in emulation on a DEC Alpha, at about the same speed as the Symbolics chips could ever run. The company still survives and is looking for its ecological niche.⁴⁶

Table 5.2: Symbolics Revenues⁴⁷



5.3 Other AI Hardware Players

In addition to LMI and Symbolics, several other firms made Lisp hardware. In particular, Texas Instruments and Xerox were active in competing for this market. TI's AI activities came from their Data Systems Group. After acquiring a 25% stake in LMI in 1983, TI decided to sell its own Lisp-chip based computers. It introduced the Explorer (built using technology from MIT and LMI) in 1984. They scored a coup in 1985 by selling \$42 million of those systems in a three-year contract to Sperry Corp., at a time when the total Lisp hardware market was only \$85 million.⁴⁸ TI also dabbled in the software side of the business, buying a 10% stake in Carnegie Group in 1985 and offering knowledge engineering consulting services starting that same year.⁴⁹ However, over time, as the Lisp market disintegrated, so did TI's

Explorer sales. "The industry as a whole did not embrace LISP-based engines," commented Les Wyatt, TI's marketing manager for the computers systems division. Early in 1991, TI discontinued the Explorer line.⁵⁰

Xerox also began selling Lisp machines from technology it had developed in its own labs. In 1985 the company introduced the 1185 (\$9995) and the 1186 (\$15865); both machines were designed to develop and run Lisp programs. Xerox's machines were at the low end of the Lisp machine market in terms of price.⁵¹ In 1988, Xerox spun out its Artificial Intelligence Business Systems Unit into a new company called Envos. The new company continued to support and sell the 1186. Unfortunately Xerox had to shut down the spin-off in 1989 and brought their products back into the Xerox fold.⁵²

Another hardware area related to AI was massively parallel computing. Firms like Kendall Square Research, formed by Henry Burkhardt in 1986, and Thinking Machines Inc., founded by Danny Hillis from the MIT AI Laboratory, built these "supercomputers" that were used for highly processor-intensive applications, such as artificial intelligence programs. These two companies, however, like much of the rest of their industry, ended up having major problems, and both firms filed for bankruptcy in 1994.⁵³ I have not focused on these companies as they were not directly part of the AI/Lisp bubble.

5.4 AI Hardware Industry Summary

Specialized hardware, which optimizes for a few functions at the expense of others, is almost always a challenging endeavor in the marketplace. While niche markets are sometimes willing to pay premiums for performance for a particular improved functionality, these markets tend to be small. And thanks to the constant improvement and price declines of "generic" multi-purpose chips, specialized chip companies must continue to improve at a rate that will keep their offerings interesting to their market. But since the multi-purpose chip companies tend to be much larger and have larger R&D budgets, specialized chip companies can find this effort to be a losing battle. In the AI field, when they found themselves losing market share, they attempted to lower prices to compete; but this action, of course, decreased their revenues and eventually made the business unprofitable.

The silver lining for specialized chip firms, however, is that radical changes in chip technology are unlikely to come from the incremental improvement of the multi-purpose chip designers. The radical restructuring mode will follow this incremental improvement, comments the MIT AI Laboratory's Tom

Knight. If a company can spearhead this leap, as many specialized chip firms attempt to do, they can reap large benefits from their effort. In the future, software implementations of prototypes may bring down the costs of chip development and may level the playing field for some of the smaller firms.⁵⁴

While management and marketing issues plagued the AI hardware industry, it also faced an uphill battle by the nature of its chosen market. The players (and their investors) took the gamble that they were going to change the face of the industry with their new technology, and thus they all would reap the benefits from doing so. Unfortunately for them, the story did not pan out that way. Today, the old "majors" that sunk the Lisp platform chips, namely Sun Microsystems, Hewlett Packard, and IBM, are facing challenges from (or have already succumbed to) the industry-standard Intel platform. Niche players like Cray still survive in small markets (supercomputing, in Cray's case), but they face a constant challenge from the multi-purpose chip companies creeping up on them from below.

¹ Tom Knight, Interview by the author, Cambridge, Massachusetts, 24 February 1999.

² Alvin Graylin, Kari Kjolaa, Jonathan Loflin, and Jimmie Walker, "Symbolics, Inc.: A Failure of Heterogeneous Engineering," (MIT Course 6.933 paper, Fall 1998).

³ Tom Knight, *Implementation of a List Processing Machine*. MIT EECS Masters Thesis, January, 1979.

⁴ Knight, 24 February 1999.

⁵ Knight, 24 February 1999.

⁶ Marc Frons, "A Semiconductor Scion Who's Making It On His Own," *Business Week* (11 November 1985), 78.

⁷ Edith Myers, "Moving Beyond Lisp; Lisp Machine Inc Has Branched Into The Worlds of Prolog and Unix," *Datamation* (15 June 1984 v30), 64.

⁸ Newquist, *The Brain Makers*, 223.

⁹ Newquist, *The Brain Makers*, 243.

¹⁰ Myers, *Datamation*.

¹¹ "Texas Instruments Acquires 25% Stake in Lisp Machine" *Dow Jones News Service* (16 August 1983).

¹² Myers, *Datamation*.

¹³ Myers, *Datamation*.

¹⁴ John Teresko, "AI Moves Into Industrial Control Arena," *Industry Week* (25 November 1985 v227), 70.

¹⁵ Newquist, *The Brain Makers*, 301-303.

¹⁶ Mary Jo Foley, "AI: The Battle Over Hardware Is Not Symbolic," *Electronic Business* (1 August 1986).

¹⁷ Newquist, *The Brain Makers*, 332.

¹⁸ Frons, *Business Week*.

¹⁹ Newquist, *The Brain Makers*, 333.

²⁰ Sources: Terry Fiedler, "The Soft Spot for AI Hardware Makers," *New England Business* (6 April 1987), 20; Newquist, *The Brain Makers*; Foley, *Electronic Business*.

²¹ Newquist, *The Brain Makers*, 343-344.

²² Foley, *Electronic Business*.

²³ Newquist, *The Brain Makers*, 346-347.

²⁴ Gensym Website, "Company Overview" (<http://www.gensym.com>)

²⁵ Newquist, *The Brain Makers*.

²⁶ Knight, 24 February 1999.

²⁷ Knight, 24 February 1999.

²⁸ Graylin, Kjolaa, Loflin and Walker.

²⁹ Knight, 24 February 1999.

-
- ³⁰ Knight, 24 February 1999.
- ³¹ Newquist, *The Brain Makers*.
- ³² Graylin, Kjolaas, Loflin and Walker.
- ³³ Knight, 24 February 1999.
- ³⁴ Newquist, *The Brain Makers*.
- ³⁵ Newquist, *The Brain Makers*.
- ³⁶ Knight, 24 February 1999.
- ³⁷ Newquist, *The Brain Makers*.
- ³⁸ Ray Weiss, "New Home for Lisp: Symbolics Goes to Sun Platform," *Electronic Engineering Times* (18 September 1989), 63.
- ³⁹ Newquist, *The Brain Makers*.
- ⁴⁰ Julie Pitta, "Where Lisp Slipped; Symbolics made the Mistake of Believing Its Own Press Clippings," *Forbes* (16 October 1989), 262.
- ⁴¹ "Lisp To Be Heard In Unix," *Unix Today!* (21 August 1989), 14.
- ⁴² Weiss, *Electronic Engineering Times*.
- ⁴³ Wendy Hower, "Change at Symbolics Must Be Deep-Rooted," *Boston Business Journal* (9 March 1992), 7.
- ⁴⁴ Newquist, *The Brain Makers*.
- ⁴⁵ Knight, 24 February 1999.
- ⁴⁶ Knight, 24 February 1999.
- ⁴⁷ Graylin, Kjolaas, Loflin and Walker.
- ⁴⁸ "The Data Systems Group of Texas Instruments Signs a Multimillion Dollar Hardware-Software Contract to Deliver Its Systems to Sperry Corp.," *PR Newswire* (24 June 1985), 86.
- ⁴⁹ "Data Systems Group of Texas Instruments Inc. Now Offering Knowledge Engineering Consulting Services," *PR Newswire* (20 November 1985).
- ⁵⁰ Kirk Ladendorf, "Tech Briefs: Austin is Home to Variety of Software Firms," *Austin American-Statesman* (7 April 1991), C2.
- ⁵¹ "Xerox Corp. Unit Unveils 2 Computer Workstations," *The Wall Street Journal* (20 August 1985).
- ⁵² "Xerox Foiled Again," *Computer System News CMP Computers File* (29 May 1989), 50.
- ⁵³ "Kendall Square Research Files for Bankruptcy," *The Associated Press* (30 December 1994).
- ⁵⁴ Knight, 24 February 1999.

6 Conclusions and Recommendations

While this thesis focuses on companies in the 1980s, there has been a large amount of activity in this field since then. Although most financiers avoided "artificial intelligence" firms in the early 1990s, several successful firms have utilized core AI technologies into their products. They may call them intelligence applications or knowledge management systems, or they may focus on the solution, such as customer relationship management, like Pegasystems, or email management, as in the case of Kana Communications.

This section will first summarize the current status of the firms profiled in this thesis. Next, I will look at a few of today's stories of companies using artificial intelligence technologies and compare their stories with those of the previous chapters. While some of these firms are still too young to qualify as lasting successes, they have shown they at least have the potential to get there.

6.1 Where Are They Now?

The current status of the firms profiled in this thesis is shown in Table 6.1 below. While a few of the companies disappeared altogether (LMI, Palladian), most survive in some form today. The legacies of a few, like Lucid and AICorp, exist as a part of a larger company, and the rest act as independent companies. The former expert systems companies, described in Table 6.1, are mostly applying their expert system technology to a particular area, such as network management or electronic customer service.

Table 6.1: Current Status of Profiled AI Firms

Name	Website	Current Status
Carnegie Group	www.cgi.com	Subsidiary of Logica plc. Consulting to call center and decision support customers
IntelliCorp	www.intellicorp.com	Business process management and application integration products and services for the ERP
Inference	www.inference.com	Software and services for online customer service; spun out Brightware (www.brightware.com), also web-based customer service, also ART*Enterprise

Teknowledge	www.teknowledge.com	Software and services for distributed knowledge processing applications
AICorp	www.platinum.com	Merged with Aion, became Trinzic; merged into Platinum Technology. Now provides business rule development environment (Platinum Aion)
Neuron Data	www.neurondata.com	Software for customer relationship management (CRM) using business rules solutions
Gold Hill	www.goldhill-inc.com	Object oriented software tools and applications
Lucid	www.harlequin.com	Bought by Harlequin, which still provides Lisp tools
Apex	www.aesclever.com	Solutions for electronic network management
Palladian		Bankrupt
Cycorp	www.cyc.com	Software and services using its common-sense knowledge base
Symbolics	www.symbolics.com	Bankrupt; assets bought by Symbolics Technology, which supports and upgrades the Genera software
LMI		Bankrupt

6.2 Modern AI Firms

I have chosen the following firms because they have similar roots (core technology and people from TI, Stanford or MIT) as the AI firms I looked at earlier in this thesis, but they have managed thus far to avoid most of the previous generation's problems. Like Symbolics and LMI, Silicon Spice is a customized chip firm from the MIT AI Lab. On the software side, Trilogy was founded by former students of Stanford's Ed Feigenbaum ("the father of expert systems"). i2 was started by a former employee of TI's AI Laboratory, and Ascent was created by the former director of MIT's AI Laboratory. All of these firms today show promise, however, in providing solutions to real problems.

6.2.1 Silicon Spice

Silicon Spice, a 1995 finalist in MIT's \$50K Entrepreneurship Competition, was founded out of the MIT AI Laboratory to create a new processing architecture for semi-conductors. They plan to

revolutionize the telecommunications industry by drastically improving the performance of telecom applications. The company is currently located in Mountain View, California.

Thus far they have been successful in raising money from top-tier venture capitalists and in bringing in experienced leaders from the semi-conductor industry.

In March 1997, Silicon Spice received first-round venture capital funding of \$3.3 million from leading venture capital firms, including New Enterprise Associates (NEA), Chemicals and Materials Enterprise Associates (CMEA) and World View Technology Partners. In April 1998, Silicon Spice closed \$7.0 million in second round funding from its first round investors with Kleiner, Perkins, Caufield & Beyers joining as a new investor. In April 1998, Vinod Dham was named President & CEO. Mr. Dham was previously Vice President and General Manager of Intel and AMD, overseeing Intel's 486 and Pentium developments and AMD's K6 development.¹

Like Symbolics and LMI, Silicon Spice is using ideas from the MIT AI Lab to build a specialized chip for use in a particular industry (in this case, telecommunications). Unlike the LISP-chip firms, however, Silicon Spice has brought in industry experts from the start (Vinod Dham is considered the "Father of the Pentium"), and has brought in enough venture capital from top VC firms to give them enough resources to fund their research and development efforts. They have also chosen an industry, telecommunications, which is large (unlike that of LISP developers) and has a lot of potential. The firm is also staying relatively quiet about their development to the press, unlike the LISP companies, and thus avoiding the more unrealistic expectations of their predecessors.²

While their long-term viability remains to be determined, Silicon Spice has at least avoided the early missteps of many tech startups, and may prove to show that specialized chip firms can break the generic chip mold.

6.2.2 i2 Technologies

i2 Technologies, the "leader in supply chain management and business optimization"³, is a modern-day expert systems company that originally focused on the manufacturing industry. Founded in 1988 as Intellection in Dallas, Texas, by Sanjiv Sidhu, a former employee of Texas Instrument's artificial intelligence laboratory,⁴ the firm believed that it could improve manufacturing planning through faster execution and focus on the business goals and the client's condition. Today i2's products cover "every phase of producing, delivering and selling goods and services." Their product family, RHYTHM, offers a solution for the needs of electronic Business Process Optimization (eBPO), through Product Lifecycle

Management, Supply Chain Management, Customer Management, Inter-Process Planning, and Strategic Planning.⁵

i2 has succeeded thus far by having an initial industry target (manufacturing); by focusing on customer value, not their own technology (in 1995, they established the goal of providing \$50 billion in value to customers by the year 2005.)⁶ Their software, which takes an object-oriented approach that allows it to solve scheduling problems by bringing the starting condition and the desired together,⁷ can be integrated with most other major industry software packages.

In December of 1994, i2 scored a coup in bringing in Greg Brady, VP of worldwide applications marketing at Oracle, to become their President of field operations.⁸ In 1995 they integrated the Carnegie Group's Caster Planning and Scheduling System (CaPSS) and Rolling Mills Scheduling System (RMSS) technologies into their Rhythm suite of applications.⁹

Unlike other AI software firms, i2 waited until the necessary computing power was available inexpensively for their applications. The power was not previously available, said Sidhu, to digitally represent the complex flow of a production process and calculate the rippling effect of a single change in that flow. They can now simulate the global impact of a change anywhere in the supply system.¹⁰

The company went public in April of 1996. It was ranked fifth on the 1999 Forbes ASAP Top 100 Technology Companies (and first in 1998) with 1998 revenues of \$362 million. i2 continues to hold promise, although its stock as of April 1999 was around 26 (giving the company a \$1.8 billion market capitalization), only 30% above its IPO price of 20.

6.2.3 Trilogy Software

Trilogy was founded by Joe Liedmandt and fellow Stanford students Christina Jones, Chris Porch and Seth Stratton and Thomas Carter in 1990. After learning about artificial intelligence technologies from Stanford AI luminaries like Ed Feigenbaum ("the father of expert systems"), and realizing that corporate sales and marketing departments were drastically under-utilizing these technologies, Liedmandt saw his window and formed his company. They combined various AI techniques such as constraint-based equations, rule-based programming, and object-oriented programming to build a system that would be a generic type of XCON for sales and marketing teams to be able to easily update and configure their product offerings for prospective customers.¹¹

In 1991, the company moved to Austin and brought in David Franke from the research consortium MCC. Franke's presence gave the firm more weight and sold their first major deal: a \$3.5 million deal with Hewlett-Packard for sales-configuration software and support services. Soon thereafter, they signed up Boeing, Silicon Graphics, and Alcatel, among others.¹²

Today, Trilogy's product lineup includes the Selling Chain suite, which contains applications that offer functionality for cataloging, commission, configuration, contracting, pricing, promotions, proposals, and quotes; and recently announced Buying Chain, which controls the corporate buying process. Trilogy also spun off PCOrder (of which Trilogy still owns 85%), which went public earlier this year and currently has a market cap of over \$1 billion. PCOrder offers e-commerce technology and services that allows computer manufacturers, distributors, resellers and end users to buy and sell products online, and uses AI technologies to determine configuration and pricing.¹³ Trilogy itself is still private and has over 800 employees.

Trilogy had many similarities to the early AI companies: students from Stanford, studying under Feigenbaum, go start an expert systems company to change the way business works. In contrast to the earlier AI expert systems companies, however, the hardware processing speed and cost was appropriate for their application. They also built their products in C++ from the start, avoiding the LISP handcuffs; they focused on particular vertical market segments, first computers, and pre-programmed a fair amount of knowledge into their system rather than offering their customers just a shell. Trilogy today also has a fairly extensive consulting operation built into their business model that helps their customers get Trilogy's products not just demonstrated but deployed.

6.2.4 Ascent Technology

Ascent Technology was founded in 1986 by Patrick Winston, Philippe Brou, Jonathan Bliss, and Karen Predergast from the MIT AI Laboratory. The company today specializes in solving

complex resource scheduling, allocation, and management problems. Our primary customer focus is on airports and airlines. We provide efficient, reliable solutions for the planning and real-time management of ground operations, such as gate management and personnel allocation, and for flight-related activities, such as aircraft routing and tracking. We combine advanced database, artificial intelligence, and data-mining technologies with industry-standard Intel, Microsoft, and Oracle platforms to provide fully integrated systems that automate resource scheduling, situation assessment, and real-time response to schedule disruptions.¹⁴

Ascent was started after the "AI bubble" as a products company, not a research firm. The founders were looking for a new kind of challenge outside of the lab. Winston focused on bringing in DARPA contracts and helping with recruiting technical people. They did not take any venture capital funding, and thus were able to grow slowly; Winston's AI Lab connection may have tainted their ability to talk to any venture capitalists, who were still scared of the AI concept.¹⁵

In the beginning, they were not sure of what product they would build. A friend from HBS suggested they get into the slowly developing resource allocation market for airlines. Previously, airlines would use magnetic boards for manual gate scheduling. Did joint development product with Continental for a gate management system. They did a joint development product with Continental Airlines in 1987 for a gate management system. It had a small expert systems component, and a huge systems integration component.

They had some difficulties selling the contract to Continental; their internal information processing subsidiaries wanted to develop it themselves, but eventually Ascent convinced the parent company. The agreement also had a clause that if either firm went bankrupt, the other firm kept the system. Supposedly this clause was because Continental worried Ascent, a startup software company, might disappear; but in fact, it was Continental who went bankrupt, and Ascent got to keep the system.¹⁶

After this system, they struggled for a few years. They tried a joint development contract with IBM's new airline system division, but nothing much came of the agreement. Then they hit the jackpot: Desert Storm. The military need to find a way to get all its troops and supplies to the Persian Gulf as quickly as possible, and Ascent's software looked to have the potential to manage the process. They sold themselves to the military, and spent ten weeks writing the DART system; they would have to go to the Air Force base in St. Louis every week from Cambridge.¹⁷

After Desert Storm, they had a reference they could use in bringing in other large clients. They signed on Delta's Atlanta hub, and started saving them \$50-100K per day. Now they have airports using their software in Hong Kong, Oslo, and Kuala Lumpur, among others. Their revenue split became about 1/3 DARPA, and 2/3 commercial. However, their sales effort, until recently, continued to struggle. Their first salesman was from DEC; however, the model at the computer firm was rather different from that of a startup software company. In 1998 they brought in Windler Schweer, a former IBM regional business manager, and the sales side of the business has blossomed. The firm today employs around 30 people.¹⁸

6.3 Lessons for High-Tech Entrepreneurship

In 1987, the Boston Globe wrote about AI Alley, the region in Cambridge's Kendall Square where many of the AI firms had set up their offices, as the home of the AI passion. "These [technologists] are no discipline problems of slack-offs [however]; they will work ferociously until dawn. In fact, they are the most skilled computer programmers around, America's hope for surpassing the Japanese. It is a subculture containing a myriad of social types: the awkward, wide-eyed hackers who send out for pizza at midnight; consultants who call Japan from their car phones; and venture capitalists looking for a stake in the future. They don't all speak the same language, but they are united by the excitement and risk of riding a largely untried technology."¹⁹

Unfortunately for most of those founders, the companies never fulfilled their original promise. They had several factors working against them, however, that the next generation of firms has hopefully avoided. First, the companies suffered from an overly academic bias in their management; this problem led to unrealistic expectations in their ability to revolutionize the corporate IT world with new hardware and software platforms and technology. Most of these firms never "crossed the chasm" by selling to the more mainstream market; their products were not ready for masses. They also underestimated the skill level of their customers and expected them to be able to put their products to use much more easily than was actually possible. Finally, many of the firms miscalculated where their products could add the most value; instead of providing expert system shells their customers could not use, they should have focused on building vertical products that integrated into their customer's current systems.

Amazingly, the hype that surrounded the AI industry let it ignore these problems for a decade before reality and the "AI Winter" set in. Perhaps without the hype to propel them along the firms would have recognized their problems sooner and focused better, but instead, by the time customers gave up on AI, it was too late for most of these companies to change. However, as many of today's high-tech companies are proving, AI does have real promise but not an industry in and of itself. But artificial intelligence technologies, added to software and hardware that does solve problems, can make those solutions even better, and may, from the inside out, revolutionize the IT world after all.

¹ Silicon Spice Website, "About Us" (<http://www.siliconspice.com>).

² Owen Thomas, "Father of Pentium to Head Silicon Spice," *Red Herring Online* (2 April 1998).

³ i2 Technologies Website, "Corporate Information" (<http://www.i2.com>).

-
- ⁴ Alan Goldstein, "Software Company Finds Its Rhythm To Be A Success," *The Dallas Morning News* (26 May 1996), 1S.
- ⁵ i2 Technologies Website, "RHYTHM Solutions" (<http://www.i2.com>).
- ⁶ i2 Technologies Website, "RHYTHM Solutions" (<http://www.i2.com>).
- ⁷ "The Uses of Time (Just-In-Time and Computer-Aided Manufacturing)," *The Economist* (5 March 1994), 15.
- ⁸ Doug Bartholomew, "Oracle VP Bolts For i2 Top Post -- Former SAP Chief Seen as Likely Successor," *InformationWeek, CMP Communications File* (19 December 1994), 24.
- ⁹ "Carnegie Group and i2 Technologies Announce First-Ever Integrated Scheduling Solution," *PR Newswire* (24 July 1995).
- ¹⁰ Amal Kumar Naj, "Manufacturing Gets a New Craze From Software: Speed," *The Wall Street Journal* (13 August 1996), B4.
- ¹¹ Josh McHugh, "'Holy Cow, No One's Done This!' (Trilogy Development Group)," *Forbes* (3 June 1996), 122.
- ¹² McHugh, *Forbes*.
- ¹³ James E. Gaskin, "PcOrder.com Boosting Web Brainpower," *Inter@ctive Week from ZDWire* (1 February 1999).
- ¹⁴ Ascent Webpage, "General Information" (<http://www.ascent.com/generalinfo.htm>).
- ¹⁵ Patrick Winston, Interview by the author (21 October 1998).
- ¹⁶ Winston.
- ¹⁷ Winston.
- ¹⁸ Winston.
- ¹⁹ Daniel Golden, "'AI Alley': A Haven for Hackers New Artificial Intelligence Industry Finds Its Natural Home in Kendall Square," *The Boston Globe* (8 March 1987), A17.

Appendix I: Bibliography

1. "A Call for Candor," *Venture Capital Journal* (1 March 1995).
2. "AICorp and Aion Sign Merger Agreement," *Business Wire* (25 June 1992).
3. "An Interview with Carl Wolf," *PC AI Magazine* (Sept/Oct 1988).
4. Ascent Website, "General Information" (<http://www.ascent.com/generalinfo.htm>).
5. Michael Ball, "Palladian & Apollo Launch 'Bargain' AI System," *Boston Business Journal* (30 March 1987), 14.
6. Michael Ball, "Ring-Around-the-Rosy: What's Behind the Recent Shake-Up?" *Boston Business Journal* (7 September 1987), 1.
7. Doug Bartholomew, "Oracle VP Bolts For i2 Top Post -- Former SAP Chief Seen as Likely Successor," *InformationWeek, CMP Communications File* (19 December 1994), 24.
8. Gordon Bell, *High-Tech Ventures* (Massachusetts: Addison-Wesley Publishing Company, 1991).
9. Pat Bellamah, "Neuron Enhances Its Expert System to Link to Many Databases," *PC Week* (25 August 1988 v5 n30), 33.
10. Kendra Berteotti, "High Hopes at Carnegie Group for New Initiative," *Pittsburgh Business Times & Journal* (30 April 1990) 2.
11. Barbara Buell, "The Professor Getting Straight As on Route 128 - Larry Harris Mothballed His Cap and Gown to Start a Company that Cashes In on Artificial Intelligence," *Business Week* (15 April 1985), 150J.
12. William Bulkeley, "Bright Outlook for Artificial Intelligence Yields to Slow Growth and Big Cutbacks," *The Wall Street Journal* (5 July 1990), B1.
13. William Bulkeley, "Computers Take On New Role As Experts in Financial Affairs," *The Wall Street Journal* (7 February 1986).
14. William Bulkeley, "Stocks of Artificial Intelligence Firms Prosper, Though Some Analysts Are Advising Wariness," *The Wall Street Journal* (31 March 1986).
15. Peter Burrows, "Artificial Intelligence Finds Work in Manufacturing," *Electronic Business Buyer* (5 February 1990), 65.
16. Martin Campbell-Kelly and William Aspray, *Computer: A History of the Information Machine* (New York: BasicBooks, 1996).
17. Ralph Emmett Carlyle, "Sneaking In the Back; Get Ready For a Quiet Influx of Expert Systems Products From Several Familiar Faces," *Datamation* (1 October 1986 v32), 32.
18. "Carnegie Group and i2 Technologies Announce First-Ever Integrated Scheduling Solution," *PR Newswire* (24 July 1995).
19. "Carnegie Group Announced the Availability of Rock 2.1," *PR Newswire* (7 July 1992).
20. "Carnegie Group Inc. Integrates Its Two Software Products, 'Knowledge Craft' and 'Language Craft'," *PR Newswire* (28 April 1986), 402.
21. "Carnegie Group Signs Contract with Caterpillar for Development of a Fully Automatic Text Translation System," *PR Newswire* (8 May 1992).
22. "Carnegie Stock Up," *Pittsburgh Post-Gazette* (5 December 1995), c-9.
23. Jean-Claude Charlet, *Firefly Network* (Graduate School of Business, Stanford University, 1998).
24. Mark Clifford, "Artificial Intelligence (Investing In High-Tech Firms)," *FW* (23 January 1985), 13.
25. Robert Clyatt, *Effective Technology Transfer Between University and Industry: A Case Study of the M.I.T. Artificial Intelligence Laboratory*, MIT Masters of Science in Management thesis (June 1985).
26. Brian Cook, "AI: Industry's New Brain Child" *Industry Week* (1 April 1991), 54.
27. Philip Cooper, *Artificial Intelligence: A Heuristic Search for Commercial and Management Science Applications*, MIT Masters of Science in Management thesis (1984).
28. Ellen T. Corliss, "Gold Hill Fights the Odds," *Mass High Tech* (30 July 1990), 1.
29. Michael Cusumano, *Japan's Software Factories* (New York: Oxford University Press, 1991).
30. "Data Systems Group of Texas Instruments Inc. Now Offering Knowledge Engineering Consulting Services," *PR Newswire* (20 November 1985).
31. "The Data Systems Group of Texas Instruments Signs a Multimillion Dollar Hardware-Software Contract to Deliver Its Systems to Sperry Corp.," *PR Newswire* (24 June 1985), 86.
32. Renee Deger, "Blinded By Science," *Venture Capital Journal* (1 March 1995).

33. John Desmond, "When MIS and AI Converge," *Software Magazine* (1 May 1988), 58.
34. Herbert Dreyfus and Stuart Dreyfus, *Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer* (New York: The Free Press, 1986).
35. Richard Duda and Edward Shortliffe, "Expert Systems Research," *Science* (15 April 1983), 261.
36. Richard Duffy, "Lisp Language for the PC Introduced; Gold Hill Says GCLisp Supports DOS, Cost \$495," *PC Week* (17 April 1984 v1), 4.
37. Esther Dyson, "Your Humble Servant (Artificial Intelligence Systems)," *Forbes* (2 November 1987), 204.
38. John Eckhouse, "Hot Investments for 1985," *The San Francisco Chronicle* (1 February 1985), 35.
39. Lawrence Edelman, "AICorp of Waltham Goes Public Sale of 2.25 Million Common Shares \$21.4 Million," *The Boston Globe* (27 June 1990), 45.
40. "Expert-system Development From IntelliCorp Now Runs on PCs," *PC Week* (5 April 1988, v5 n14), 26.
41. "The Fall of Supercomputers," *Boston Herald* (16 August 1994), 27.
42. Edward Feigenbaum and Pamela McCorduck, *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World* (2nd edition, New York: Addison-Wesley Publishing Company, 1984).
43. Marc Ferranti, "Neuron readies more flexible Nexpert," *PC Week* (2 July 1990 v7 n26), 43.
44. Mary Jo Foley, "AI: The Battle Over Hardware Is Not Symbolic," *Electronic Business* (1 August 1986).
45. "Ford Acquires 10% Stake in L.A. Software Company," *Los Angeles Times* (25 October 1985), 2.
46. "Ford Motor Co. Acquires 10% Interest in Carnegie Group Inc. and Inference Corp.," *PR Newswire* (24 October 1985).
47. Franz Website, "About Franz" (<http://www.franz.com>).
48. Otto Friedrich, "The Computer Moves In," *Time* (3 January 1983), 14.
49. Marc Frons, "A Semiconductor Scion Who's Making It On His Own," *Business Week* (11 November 1985), 78.
50. Joyce Gannon, "Perq Lays Off 30; Spins Off AI Firm," *Pittsburgh Business Times & Journal* (28 January 1995), 1.
51. Joyce Gannon, "Start-Up Slowed: Software Firm Says Loan Dispute Caused Bankruptcy Filing," *Pittsburgh Post-Gazette* (19 September 1991), 14.
52. James E. Gaskin, "PcOrder.com Boosting Web Brainpower," *Inter@ctive Week from ZDWire* (1 February 1999).
53. Gensym Website, "Company Overview" (<http://www.gensym.com>).
54. Paul Gillin, "Plans Dropped for PC Version of 'Intellect,'" *PC Week* (5 March 1985 v2), 123.
55. Paul Gillin, "LISP Packages Bespeak Success for AI Tools Firm," *PC Week* (1 April 1986 v3), 101.
56. Gold Hill Website, Background Information, <http://www.goldhill-inc.com>
57. Daniel Golden, "'AI Alley': A Haven for Hackers New Artificial Intelligence Industry Finds Its Natural Home in Kendall Square," *The Boston Globe* (8 March 1987), A17.
58. Alan Goldstein, "Software Company Finds Its Rhythm To Be A Success," *The Dallas Morning News* (26 May 1996), 1S.
59. Alvin Graylin, Kari Kjolaas, Jonathan Loflin, and Jimmie Walker, *Symbolics, Inc.: A Failure of Heterogeneous Engineering*, (MIT Course 6.933 paper, Fall 1998).
60. Udayan Gupta, "Concerns Rush To Go Public As Technology Issues Prosper," *The Wall Street Journal* (24 March 1986).
61. John Halusha, "General Motors Corp, in Little-Noticed Move, Pays \$3 Million for 11 Percent Interest in Teknowledge Inc.," *New York Times Abstracts* (3 May 1984), 1.
62. Harlequin Website, "About Harlequin" (<http://www.harlequin.com/about>).
63. Wendy Hower, "Change at Symbolics Must Be Deep-Rooted," *Boston Business Journal* (9 March 1992), 7.
64. i2 Technologies Website, "Corporate Information" (<http://www.i2.com>).
65. i2 Technologies Website, "RHYTHM Solutions" (<http://www.i2.com>).
66. "Inference Closes \$5.5 Million in Growth Financing; New Round of Funding Led by JP Morgan," *Business Wire* (15 July 1991).
67. IntelliCorp SEC Annual Report (Form 10-KSB), For Fiscal Year Ended June 30, 1998.
68. "IntelliCorp Sells Remaining Interest in Joint Venture to Amoco," *Business Wire* (9 November 1990).

69. "IntelliGenetics' 1.6 Million Initial Shares Offered at \$6," *Dow Jones News Service - Ticker* (15 December 1983).
70. "Interview: Jo Marks," *PC AI Magazine* (Jan/Feb 1995), 39.
71. Jan Johnson, "Easy Does It: Natural Language Query Systems Have Gained New Attention in Recent Months," *Datamation* (15 June 1984 v30), 48.
72. R. Colin Johnson, "Layoffs Hit AI Specialist Gold Hill," *Electronic Engineering Times*, CMP Electronics File (12 March 1990), 15.
73. 8. Jerry Kaplan, *Startup* (New York: Penguin Books, 1994).
74. Kevin Kelly, "A High-Tech Think Tank Thinks Big Bucks," *Business Week* (25 September 1989), 222.
75. "Kendall Square Research Files for Bankruptcy," *The Associated Press* (30 December 1994).
76. Barbara Klein, "A Risk-Taker Sizes Up the Rewards," *Pennsylvania Business and Technology* (1 April 1992) 32.
77. Tom Knight, *Implementation of a List Processing Machine*. MIT EECS Masters Thesis, January, 1979.
78. Karen Kovatch, "Carnegie Group Seeks New Clients to Bolster Investors," *Pittsburgh Business Times & Journal* (16 June 1997), 5.
79. Karen Kovatch, "Unfulfilled Promise?" *Pittsburgh Business Times & Journal* (23 October 1998), 1.
80. Kirk Ladendorf, "Tech Briefs: Austin is Home to Variety of Software Firms," *Austin American-Statesman* (7 April 1991), C2.
81. "Lisp To Be Heard In Unix," *Unix Today!* (21 August 1989), 14.
82. "Lockheed to Assume Minority Position in Inference Corp.," *Dow Jones News Service - Ticker* (21 December 1984).
83. "Lockheed to Invest Additional \$2 Million in Inference Corp.," *Dow Jones News Service - Ticker* (4 March 1986).
84. "Logica to Acquire Carnegie Group for \$35 Million," Logica plc Press Release (1 October 1998).
85. Seth H. Lubove, "Carnegie Group President's Resignation Shows Tension at Some Start-Up Firms", *The Wall Street Journal* (20 February 1987).
- "Lucid launches Advanced C++ and C Programming System for SPARC; Provides a Natural Path to Object-Oriented Programming," *Business Wire* (2 January 1992).
86. Josh McHugh, "'Holy Cow, No One's Done This!' (Trilogy Development Group)," *Forbes* (3 June 1996), 122.
87. Microsoft Encarta Encyclopedia, Web Edition, "Gang of Four."
88. Jeff Moad, "Building a Bridge to Expert Systems; IntelliCorp Claims It's Found a Way to Link Its Knowledge Systems with Relational Databases," *Datamation* (1 January 1987 v33), 17.
89. Geoffrey Moore, *Crossing the Chasm* (HarperBusiness 1991), 22-23.
90. Paula Musich, "Software Upgrade, AI Program Boost AS/400's PC Connectivity," *PC Week* (11 September 1989), 6.
91. Edith Myers, "Moving Beyond Lisp; Lisp Machine Inc Has Branched Into The Worlds of Prolog and Unix," *Datamation* (15 June 1984 v30), 64.
92. Amal Kumar Naj, "Manufacturing Gets a New Craze From Software: Speed," *The Wall Street Journal* (13 August 1996), B4.
93. "Neuron Data Inc (announced support for Microsoft Corp.'s Windows NT operating system)," *InfoWorld* (23 November 1992 v24 n47), 15.
94. Neuron Data Website, "Company Information" (<http://www.neurondata.com/Company/index.html>).
95. Harvey Newquist, *The Brain Makers: Genius, Ego and Greed in the Quest for Machines that Think* (Indianapolis: Sams Publishing, 1994).
96. Peter Norvig, *Paradigms of Artificial Intelligence Programming: Case Studies in Common Lisp* (<http://www.norvig.com/paip-preface.html#whylisp>).
97. Brian Oakley and Kenneth Owen, *Alvey: Britain's Strategic Computing Initiative* (Cambridge, MA: MIT Press, 1989).
98. "Oxford Molecular Acquires IntelliGenetics, Inc." *PR Newswire* (25 August 1994).
99. John Pallatto, "Expert-system Shells Span DOS, OS/2 Nets," *PC Week* (28 August 1989 v6 n34), 6.
100. John Parker, "Merger Reflects New Focus for Expert Systems," *PC Week* (2 April 1990 v7 n13), 116.
101. William Pat Patterson, "Natural language lets anyone use computers," *Industry Week* (2 April 1984 v221), 62.
102. PCAI Interview: Patrick Perez, *PC AI Magazine* (Nov/Dec 1989), 34.

103. Julie Pitta, "Where Lisp Slipped; Symbolics made the Mistake of Believing Its Own Press Clippings," *Forbes* (16 October 1989), 262.
104. "PLATINUM technology and Trinzic Complete Merger," Platinum News Release (25 August 1995).
105. Edward Roberts, *Entrepreneurs In High Technology: Lessons from MIT and Beyond* (Oxford University Press 1991).
106. Ronald Rosenberg, "Palladian Software's President Resigns," *The Boston Globe* (27 August 1987), 57.
107. Daniel Rosenheim, "Silicon Valley Slump -- It's Not All That Bad," *The San Francisco Chronicle* (23 August 1985), 6.
108. Matt Rothman and Emily Smith, "The Leading Edge of 'White-Collar Robotics' - These Hot Startups Are Rushing to Cash In On Computer Software That Mimics Human Reasoning," *Business Week*, (10 February 1986), 94.
109. Paul Sagawa, "The Carnegie Group" (Harvard Business School case study; originally published 1990; revised 11/10/92; #9-690-033).
110. Martha Scarf Samuelson, *The Start-Up of an Artificial Intelligence Company*, MIT Master of Science in Management thesis (June 1986).
111. Jim Schembari, "Firm Positioned to Gain From New Technology," *Chicago Sun-Times* (17 November 1985), 4.
112. Sy Schoen and Wendell Sykes, *Putting Artificial Intelligence to Work: Evaluating and Implementing Business Applications*, (New York: John Wiley & Sons, 1987).
113. Silicon Spice Website, "About Us" (<http://www.siliconspice.com>).
114. Emily Smith, "A High-Tech Market That's Not Feeling The Pinch – Eager Investors Have Created a Boom in Artificial Intelligence, But Can All These Startups Deliver?" *Business Week* (1 July 1985), 78.
115. Robert Sonenclar, "Going Public; The New Issues of 1986," *FW* (18 March 1986), 20.
116. Gregory Spector, "Interface Helps Build Knowledge-Based Systems," *PC Week* (29 April 1986 v3), 18.
117. "Sperry Corp. Signs Contract to Purchase Texas Instruments' Explorer Systems as Part of Its Knowledge System Product Line," *PR Newswire* (25 June 1985), 66.
118. John Sviokla, *PlanPower, XCON, and MUDMAN*, vii.
119. John Sviokla, *Building the Knowledge-Based Organization* (Harvard Business School Publishing, 1990).
120. John Sviokla, *Business Implications of Knowledge-Based Systems* (Harvard Business School Publishing, 1986).
121. John Sviokla, *An Examination of the Impact of Expert Systems on the Firm: The Case of XCON* (Harvard Business School Publishing, 1990).
122. John Sviokla, *Expert Systems Project Management Module* (Harvard Business School Publishing, 1990).
123. John Sviokla, *Expert Systems Technology and Its Transfer* (Harvard Business School Publishing, 1990).
124. John Sviokla, *Knowledge Workers and Radically New Technology* (Harvard Business School Publishing, 1996).
125. John Sviokla, *PlanPower, XCON, and MUDMAN: An In-Depth Analysis into Three Commercial Expert Systems in Use* (Harvard Business School Doctoral Thesis, 1986).
126. John Sviokla, *Turning Data into Dollars* (Harvard Business School Publishing, 1996).
127. "Systems Integration Report: Expert System Cleverly Diagnoses SNA Connectivity Problems," *Manufacturing Automation* (1 December 1994).
128. Darryl K Taft, "Art-ADA Marries Expert Systems, Ada Technology," *Government Computer News* (8 January 1990), 33.
129. Teknowledge Website, "Corporate Information" (<http://radon.teknowledge.com/overview/index.html>).
130. John Teresko, "AI Moves Into Industrial Control Arena," *Industry Week* (25 November 1985 v227), 70.
131. "Texas Instruments Acquires 25% Stake in Lisp Machine" *Dow Jones News Service* (16 August 1983).
132. "The Future of Technology in the Financial Services Industry," *American Banker* (14 April 1986), 14.
133. "The Uses of Time (Just-In-Time and Computer-Aided Manufacturing)," *The Economist* (5 March 1994), 15.
134. Owen Thomas, "Father of Pentium to Head Silicon Spice," *Red Herring Online* (2 April 1998).

135. J. Marshall Unger, *The Fifth Generation Fallacy: Why Japan Is Betting Its Future on Artificial Intelligence* (New York: Oxford University Press, 1987).
136. Ray Weiss, "New Home for Lisp: Symbolics Goes to Sun Platform," *Electronic Engineering Times* (18 September 1989), 63.
137. Patrick Winston and Karen Prendergast, *The AI Business: The Commercial Uses of Artificial Intelligence* (Cambridge, MA: MIT Press, 1984).
138. "Xerox Corp. Unit Unveils 2 Computer Workstations," *The Wall Street Journal* (20 August 1985).
139. "Xerox Foiled Again," *Computer System News CMP Computers File* (29 May 1989), 50.

Appendix II: List of Interviews

1. Bill Mark (telephone), CA, 4 February 1999.
2. Bob Fondiller (telephone), CA, 8 February 1999.
3. Chuck Williams (telephone), CA, 23 February 1999.
4. Don Barker (telephone), WA, 11 January 1999.
5. Doug Keith, Boston, MA, 21 January 1999.
6. Hubert Dreyfus, Berkeley, CA, 30 December 1998.
7. Edward Feigenbaum, Stanford, CA, 7 January 1999.
8. Fred Luconi (telephone), FL, 10 February 1999.
9. Harvey Newquist (telephone), NY, 8 January 1999.
10. John Price, Austin, TX, 9 March 1999.
11. Judy Bolger (email), 9 February 1999.
12. Ken Ross, Mountain View, CA, 13 January 1999.
13. Marvin Minsky (email), 17 February 1999.
14. Philip Cooper (telephone), New York, NY, 19 November 1998.
15. Pushpinder Singh, Cambridge, MA, 18 November 1998.
16. Qunio Takashima (email), 3 December 1998.
17. Ramanathan Guha (email), 24 February 1998
18. Randy Davis, Cambridge, MA, 20 October 1998.
19. Ed Roberts, Cambridge, MA, 11/3/98, 11/23/98, multiple 1998-9.
20. Russ Siegelman (telephone), CA, 7 April 1999.
21. Tom Knight, Cambridge, MA, 24 February 1999.
22. Patrick Winston 10/21/98, Cambridge, MA, multiple 1998-9.

I also had informal conversations with Ken Morse, Joe Liemandt, Joost Bonsen, Michael Cusumano, and Yezdi Lashkari.