

The Industrialization of Artificial Intelligence: From By-Line to Bottom Line

S. Jerrold Kaplan

*Teknowledge, Inc.
Palo Alto, CA 94301*

Over the past few years, the character of the AI community has changed. AI researchers used to be able to go about their work in peace, while the rest of the world ignored them. As the promise of practical applications of AI has slowly become reality, new players have entered the field, changing its nature forever. The quiet, intellectual community of AI researchers has been augmented by a hoard of other interested parties, including the press, the financial community, and the technology entrepreneurs.

Since we cannot go back and hide in the ivory tower, we may as well take the time to explore our new environment. I invite you to join me in a guided tour of the new AI community. Let's begin by analyzing the basic motivation for commercial interest in AI.

The "Value-Added" of Artificial Intelligence

Artificial intelligence is not in itself a commercial field, but a science and a technology. As an academic discipline, it is a collection of concepts and ideas which are appropriate for research, but which cannot be marketed. However, artificial intelligence is the scientific foundation for several growing commercial technologies. The three most important are robotics and vision (which for simplicity's sake I will consider as one field), natural language processing, and knowledge engineering. Supporting markets are forming which will

provide hardware and software tools for use in these three areas.

Artificial Intelligence vs. Natural Intelligence

The potential value of Artificial Intelligence can be better understood by contrasting it with natural intelligence. Artificial intelligence has several important commercial advantages:

- AI has *permanence*. Natural intelligence is perishable from a commercial standpoint in that employees can change their place of employment or forget information and techniques. Artificial intelligence, however, is as permanent as computer systems and programs.
- AI offers *ease of duplication*. Transferring a body of knowledge from one person to another usually requires a lengthy process of apprenticeship; even so, expertise can never be naturally duplicated. However, when knowledge is embodied in a computer system, it can be copied and easily moved to another location.
- AI can be *less expensive* than natural intelligence. There are many circumstances in which buying a small or even a large computer costs less than having corresponding human power carry out those same tasks.
- AI—being a computer technology—is *consistent and thorough*. Natural intelligence is erratic because people are erratic; they don't perform consistently.

- AI is *documentable* Decisions made by a computer can be documented by tracing the activities of the system. Natural intelligence is difficult to reproduce in the sense that a person may reach a conclusion, but at some later date be unable to recreate the reasoning process that led to it or even to recall the assumptions that were a part of it.

Let me stress that I am in no way suggesting that artificial intelligence is generally superior to natural intelligence—simply that it is different. To present a balanced picture, let's take a brief look at some advantages of natural intelligence:

- Natural intelligence is creative, whereas artificial intelligence is rather uninspired. The ability to acquire knowledge is inherent in human beings, but, with artificial intelligence, tailored knowledge must be built into a carefully constructed system
- Natural intelligence enables people to benefit by and use sensory experience directly, whereas artificial intelligence systems must work with symbolic input.
- Perhaps most important of all, human beings in their reasoning processes are able to make use at all times of a very wide context of experience and bring that to bear on each individual problem, whereas artificial intelligence systems typically gain their power by having a very narrow focus.

A look, however, at the previously mentioned advantages of AI over natural intelligence—permanence, ease of duplication, cost effectiveness, consistency and thoroughness, and documentability—reveals that the list really describes the advantages of computer systems in general, not just AI systems. This seems to suggest that the commercial advantages of AI are derived from the nature of computational systems themselves, not from symbolic processing techniques. I'd like to explore further this intriguing hypothesis by contrasting two fields which I see as structurally related: data processing and knowledge engineering.

Knowledge engineering is the representation and use of symbolic knowledge in electronic form to solve problems that normally require human intelligence or attention. The primary difference between data processing and knowledge engineering is that data processing deals with the representation and use of data, whereas knowledge engineering deals with the representation and use of knowledge. Data processing is concerned with algorithmic, repetitive processes; knowledge engineering is concerned with judgmental and inferential processes. Data processing is focused on the efficient handling of large data bases and large numbers of transactions, but knowledge engineering is concerned with the effective handling of knowledge bases and decisions. Despite the differences, there is a clear correspondence between the goals of these two fields.

In short, a hard look in a broader context reveals that AI simply extends the current commercial advantages of computer systems—their speed, accuracy, consistency, availability, and affordability—from clerical to intellectual applications.

One final thought on this subject: This extension of the role of computer systems is likely to go essentially unnoticed by the end-users of these systems. Most people interact with computers in limited ways—they do not have the same sense of the boundaries of computer technology that a computer scientist develops. To most people, computers are tools that help them solve certain specific problems—word processing, filing, billing, budget planning, etc. Each year, new and more useful applications are added. AI-based systems will simply expand the range of new applications most people experience—its business as usual. Indeed, good evidence for this is that most companies incorporating AI end-user systems into their products (as opposed to selling tools to the AI community) rarely mention AI in their promotional material.

The Development of Commercial Markets

I am often asked by people investigating the field what industries AI will benefit. Where will it be used—oil, medicine, manufacturing, consumer goods, etc.? The problem with this question is that it is a little like asking where computers can be used. The benefits do not fall along normal industry lines, but rather along certain classes of use.

Nonetheless, our experience at Teknowledge to date does suggest that there are three basic segments to the commercial market for knowledge systems (the end-product of knowledge engineering projects). These are the industrial segment, the service segment, and the management segment.

Some typical industrial applications are in scheduling, production planning, maintenance, process control, and quality assurance. Many of the systems you are likely to have heard about, in the oil industry, mining, manufacturing, etc. are likely to fall into this category.

In the service sector, some typical applications are in automation of routine advice, standardization of service, quality control, worker productivity, advice on accessing and using computer systems, and training. Perhaps the most interesting of these applications are the first two, automation of routine advice and standardization of service. Knowledge-based systems can handle relatively routine questions that require a minimum of intelligence. For example, a person opening a bank account usually needs advice on what type of account is most appropriate, given his or her particular mix of financial characteristics. Given the changing character of the banking industry, systems providing routine advice that accurately reflects the particular services of that bank would have a tremendous impact on, and value to, that organization. Second, in terms of standardizing service, it is very difficult for banks to ensure that their personnel are trained in the same ways, give the same advice, and provide the same level of service. Using a knowledge-based system would ensure that the level of service that is provided to each customer is constant, controlled, and monitored. These same advantages accrue to a wide variety of service industries that are concerned with the dissemination of accurate and timely

knowledge, such as travel planning, real estate, accounting, medicine, and sales consulting.

Surprisingly, the third segment of the market for knowledge systems is the management area. One important problem, in fact the single most critical problem for today's large corporations, is the syndication of "corporate culture." Each corporation has its own view of what it is, how it does business, how it is positioned. It is extremely difficult to inculcate an entire organization—which may encompass thousands of people—with the corporate culture. A system which distributes the corporate culture by providing advice and guidance to people faced with business decisions could help to focus and align the efforts of the organization. Knowledge systems could also be built to aid in strategic planning and fusing information from many diverse sources (where managers may not have the time or inclination to cull out what they really need)

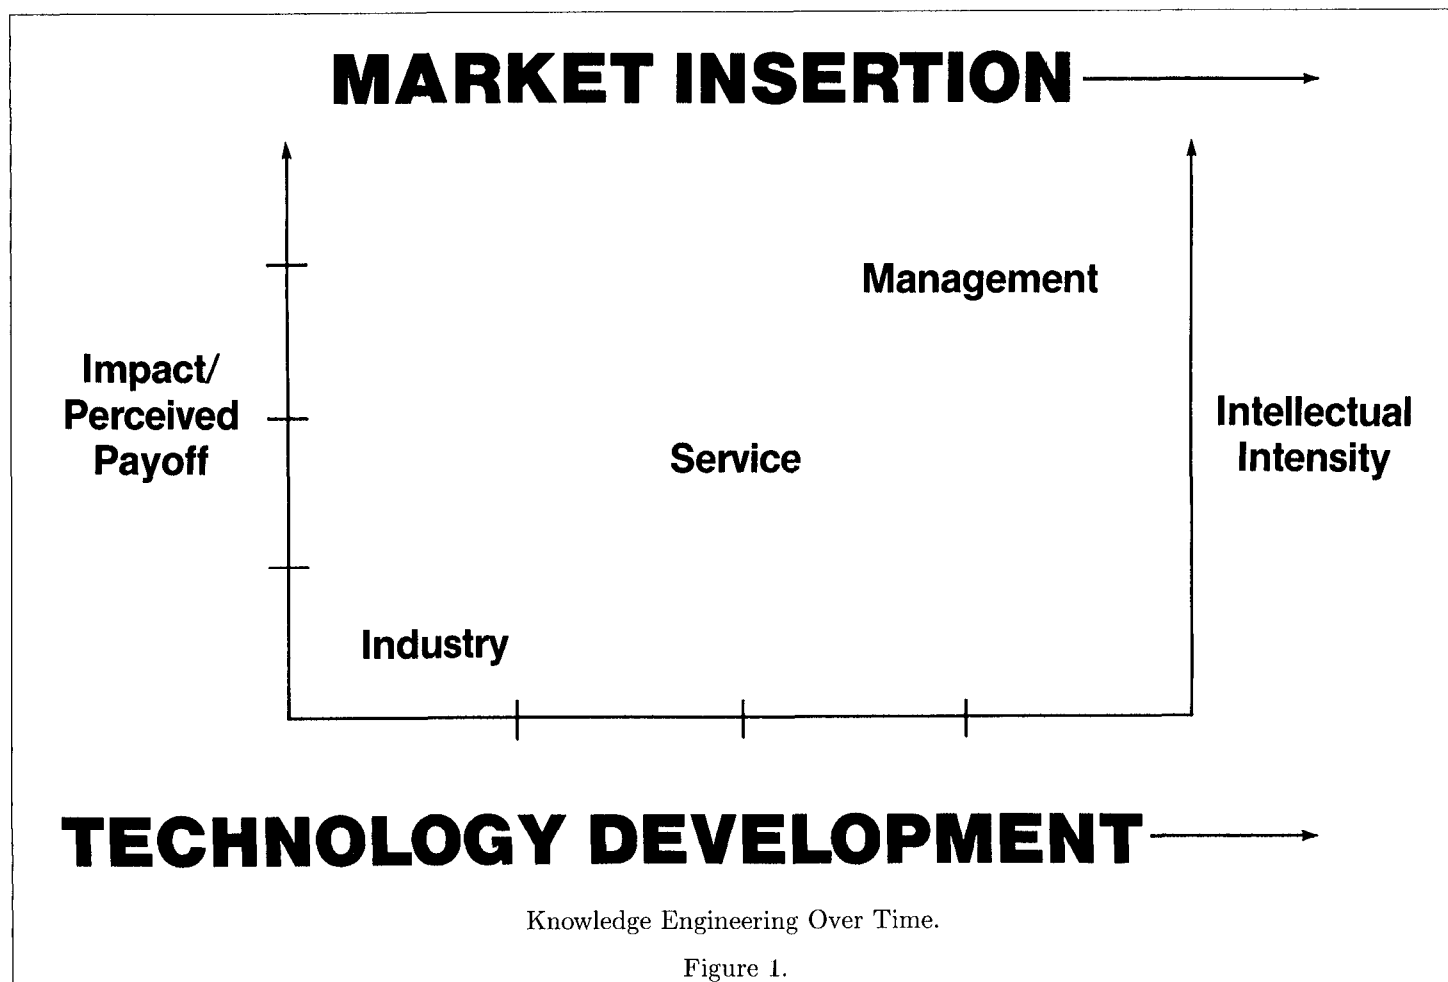
A structure is emerging for the insertion of knowledge-based technology into the industrial, service, and management market segments. Figure 1 is a graph of the development of knowledge engineering technology over time versus the payoff or perceived impact of the systems.

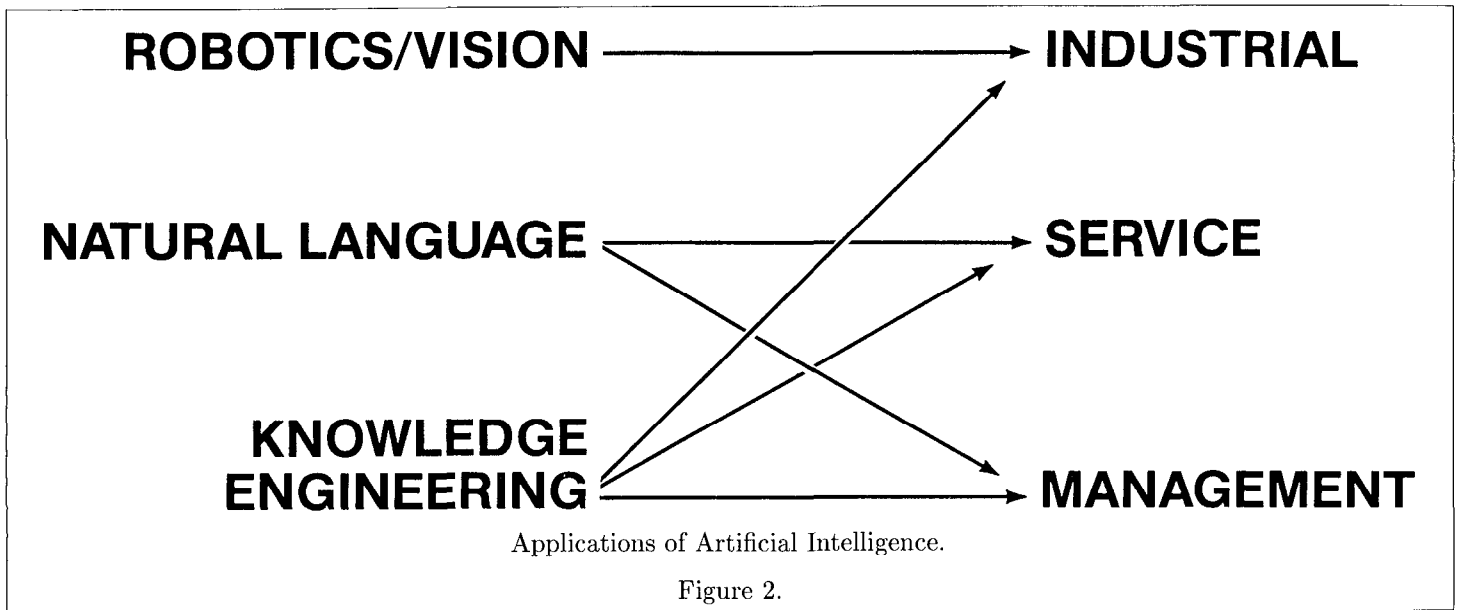
For a variety of reasons, principally social, the insertion of knowledge engineering technology into the commercial

sector appears to be moving from industrial, to service, and finally to management applications

Industrial applications have been the first major area of impact, primarily because these potential applications are the most visible and most accessible to knowledge engineers. Now, we are seeing increased activity in the service sector, as people in that sector—less accustomed to seeking out technological innovations—begin to recognize the potential value of knowledge engineering to their businesses. Finally, as central corporate staffs become aware of the technology, they are considering applications to the management of organizations. There has been little actual work and even less publicity about this last area, but it is growing. Although industrial applications have been explored further than management applications, the payoffs have been seen and measured in these systems. This results in an increased intellectual intensity in activity. Thus, the payoffs of operational systems will escalate as the focus shifts from industry to service to management.

In a sense, all commercial problems are really management problems: they result from failure to plan and apply resources effectively, misunderstanding of markets, and inability to recognize and exploit organizational strengths to achieve tactical advantages. We may eventually find that





the highest payoffs will be derived from applications in the management area. Improving the productivity of a manufacturing facility is important, but not as important as correctly determining whether the plant should be operated in the first place.

Turning to the other various commercial technologies within AI, we note that each major area—robotics and vision, natural language, and knowledge engineering—has a natural market segment which it addresses. Figure 2 points out that robotics and vision mainly have applications in the industrial segment, natural language mostly deals with the service and management segments, and knowledge engineering can be applied to all three.

The Barriers to Commercialization

Although the press glorifies the promise of AI, there are some significant barriers to the commercialization of AI technologies.

The first and biggest barrier is the abstract quality of artificial intelligence. This makes AI difficult for customers to comprehend and distinguish from other technologies. Many people who perceive a great value in artificial intelligence do not know how to apply it. In fact, there is a growing subindustry of people who help organizations decide where to apply artificial intelligence. Artificial intelligence could be described as a solution in search of a problem.

A second barrier to commercialization is unrealistic expectations about the actual capabilities of the technology and the length of its gestation period. When people don't see results occur as soon as they expect, they may be disappointed.

There are three basic sources of unrealistic expectations in today's arena. The first is the amount of media attention given to the field. There are a great many articles published about artificial intelligence; a kind of "AI fever" is sweeping

the press. Unfortunately, if the media don't see results that they can say something interesting about within the next year, they may swing the other way.


The second source is charlatanism, most of which is unintentional. Just about everybody is climbing on the AI bandwagon. Many companies are falsely claiming that their products "contain AI," as though this were an ingredient they recently added.

A major source of unrealistic expectations is being created by excessive government funding commitments combined with a shortage of qualified AI researchers. In the next few years, the government will probably submit hundreds of millions of dollars' worth of requests for proposals involving AI research. If there are not enough researchers to carry out these projects, government funding for AI may be cut later on.

Another barrier to the commercialization of artificial intelligence is the slow growth of the technical prerequisites and development of expertise. While there is a distinct shortage of experienced personnel who have a background in artificial intelligence, the people who do have that experience often do not have an understanding of real business problems. Many of the new AI companies must reeducate their academically oriented technicians to understand what value the clients will receive from the systems that the company develops.

Besides a shortage of human expertise, there is also a shortage of reliable, supported tools for performing AI activities. Many companies, including Teknowledge, are working to make the technology more accessible by supplying more sophisticated commercial tools.

Curiously, a significant barrier to commercialization can be traced directly to the terms "artificial intelligence" and "expert systems" themselves, with their cognitive connotations. If we think of artificial intelligence as simply imitating natural intelligence, we will miss its real commercial value. This "machine replaces human" orientation obscures the true



potential of the field—extending the use of the computer as a commercial tool. Accepting and understanding the technology would be much easier if the field were called something like “symbolic processing”—arguably a much more descriptive name. To compound this problem, the term “artificial intelligence” has a rather bad reputation in some circles—it is identified with large, slow, impractical systems.

The term “expert systems” is equally problematic. It suggests that the role of the technology is to replace experts. In fact, knowledge engineering is a means to capture, reason about, distribute, and use symbolic knowledge in an organization. This knowledge doesn’t necessarily have to come from experts; in fact, for practical reasons it is greatly preferable if it doesn’t. It may come from manuals, documents, databases, procedures, schedules, or other passive sources, but it becomes much more useful when represented and distributed electronically in a symbolic form. At Teknowledge, when we are asked to identify high-value applications, we seek out “knowledge bottlenecks,” where knowledge is poorly distributed, rapidly changing, or too voluminous to handle conveniently on a routine basis. Most often this does not involve an “expert.”

Finally, utilizing AI technology within a commercial organization requires imagination and creativity. AI systems may change the structure of the business, rather than augment or replace existing systems. The application of AI will create novel opportunities, but actually identifying these opportunities might require some creative thinking..

The Impact of Commercialization on AI Researchers

Until recently, AI researchers have worked mostly in the

academic community. As they have ventured out into the commercial sector, they have discovered that the rules of the game are different. Techniques useful for success in academic institutions do not transfer well to business organizations.

For instance, consider the problem of selecting an appropriate domain of application for a system. In the academic community, research must lead to publishable papers. Therefore, selected applications should be interesting to peers and resistant to criticism (“That’s an easy problem”). It is often helpful to select a problem requiring collaboration with professors in other departments—they may prove valuable allies when applying for grants, or being considered for tenure. Consequently, the classic academic applications have been extremely sophisticated problems, such as medical diagnosis and scientific discovery. By picking difficult problems, the research has contributed tremendously to our understanding of the field.

In a commercial environment, it is best to select easier problems. These may be less interesting, but they usually have payoffs comparable to those of hard problems and lower technical risk. Applications are selected to improve productivity, enhance products and services, and increase profits.

Most academic researchers develop only prototype systems. However, selecting a problem and implementing a prototype are only a very small portion of the overall effort required to field and maintain an operational system. A commercial AI company must be able to handle documentation, training, and quality assurance; analyze the requirements for a final, fielded system; and integrate, operate, and maintain that system. Figure 3 shows that these activities absorb most of a company’s time and resources.

Finally, artificial intelligence researchers must undergo

an unusual psychological transition when they move from the academic to the business community. In the academic world, researchers compete internally with their colleagues for resources, students, and tenure slots. However, they tend to cooperate externally, in part, because funding agencies favor inter-institution or inter-disciplinary proposals. In a business environment, researchers cooperate internally to achieve results, but compete externally with other companies for a share of the market and for commercial and technical advantage.

A Look at the Future

Venturing into the dangerous province of the futurist, let me make some modest predictions on the development of the field for the next few years, followed by some wild speculation on where we might be after the next few decades.

One question is whether artificial intelligence will remain a distinct discipline, carried out by identifiable and separate groups, or whether it will be subsumed into the larger mass of data processing and computer technology in general. It now seems clear that separate artificial intelligence groups will perform artificial intelligence work. For example, many large corporations, probably 40 or 50 in the U.S. alone, have formed knowledge engineering departments and are providing their staffs with equipment separate from the company's data processing equipment. This implies that, in the short run at least, artificial intelligence will remain a distinct technology.

A second question is how the market for artificial intelligence products will be structured. The market now forming can be described as three hierarchical levels. At the bottom are AI tool suppliers, who supply hardware and software to development groups. These groups—the next hierarchical level—use these tools to build systems for specific applications, so-called “vertical systems.” The top level are the purchasers of these problem-solving systems. As discussed earlier, the fact that they are buying artificial intelligence systems may be invisible to this market.

Let's look in more detail at the suppliers of AI tools. There are two types of tools—hardware tools and software tools. On the hardware side, virtually every major computer manufacturer is watching the market closely. Since the barriers to entering the hardware tools market are relatively low, in the next couple of years there will probably be heavy competition in the manufacture of work stations for symbolic processing.

Similarly, there is likely to be a flood of software tools entering the market. For a period of time, these tools will be nearly indistinguishable, as similar claims will be made for the best and the worst of them. Unlike the benchmarks for hardware, however, there is no easy way to compare the power and quality of software tools. Compounding this problem, most first-time tool purchasers don't know how to use the tools, and, therefore, can't easily determine if one is more appropriate for their problem than another. Over time,

however, the market will shake out. Software tool suppliers that can exploit practical experience to design useful tools, help train their customers to select problems, and provide applications consulting will dominate in this market.

Vertical systems will probably be built by two different types of organizations. The first type consists of internal departments of major corporations that for all practical purposes function as specialized, captive companies. These will use AI tools to enhance and deliver the products and services of their parent organization.

The second type consists of specialized companies which usually attack some particular vertical market segment. These organizations are not “selling AI” rather, they are selling a product or service which incorporates AI to some other market. For example, an AI company that builds expert systems for financial institutions is in the financial consulting business, and uses AI as a means to deliver their service. In short, the organizations that build vertical systems are quite distinct from the organizations that supply them with the tools they need.

Looking into the crystal ball, it is interesting to speculate on where we might find things in the more distant future. I will discuss the three major commercial AI technologies separately.

The development of robotics and vision systems will probably lead to the creation of much more generalized machines. Today, a wide variety of machines provides specific solutions to specific kinds of problems. In the future, there will be more general machines, which will provide a cost-effective replacement for a number of more specialized machines. As with general computing machines, volume production will offset the individual inefficiency of applying these machines to individual tasks for which they may be less than optimal.

To bring this point down to earth, imagine the effect on the typical household of a robot that could perform general chores. If this device were to exist, it would provide an inexhaustible pool of cheap domestic labor, and the economic structure for building special-purpose machines would disappear. Dishes could be washed by hand (so to speak), as could clothes, eliminating the dishwasher, washer, and dryer. Cans would be opened “manually.” Vacuums, mops, carpet sweepers, brooms, and other labor-saving cleaning aids would be unnecessary, since surfaces could simply be scrubbed or brushed when required. In this vision, homes in the next century would look more like those in the last century than those of today—simpler, more comprehensible, less technological.

Natural language will eventually provide the means for people to interface effectively with the remaining simplified technology, and to increase its accessibility and utilization. Today, most people are somewhat alienated from the technology around them. With the development of robotics and natural language, people will feel much more integrated; they will have a greater sense of control over their environment.

Eventually, knowledge engineering systems will perform

simple, day-to-day cognitive tasks, such as giving directions, planning, aiding human memory, directing traffic, scanning the news to select information of interest, helping to select merchandise and vendors. More important, knowledge engineering will revolutionize the way knowledge is captured and distributed. Today, there are two ways to capture and distribute knowledge—through teaching and through books, films, and other media. However, these methods and the knowledge they transmit are rarely as accessible or useful as a knowledge system could be.

One interesting effect will be to improve communication and flow of information in our society. Since markets (such as stock exchanges) thrive on the differential distribution of information, we can expect these functions to abate, or disappear altogether. There is no reason to advertise or promote if every potential customer has access to perfect information about the suitability of a given product or service for their specific needs. In this case, nearly everything would become a commodity.

Besides improving the overall quality of our lives, these changes are likely to alter many people's perception that they are drowning in masses of useless information. Knowledge will be available where and when it is needed, in a useful, integrated form. Rather than being inundated with radio, television, magazines, junk mail, and advertising, life should seem simpler and quieter.

Conclusion

As AI technologies have begun to fulfill their commercial promise, new constituencies have appeared with vested interests in developing and applying artificial intelligence outside

of the research environment. These newcomers have created some concern or discomfort in the academic community. How will this expanded interest affect AI researchers and their institutions? Will it disrupt the flow of students, faculty, and funding? Will the AAAI annual meetings turn into massive trade-show extravaganzas?

The first step to answering these questions lies in recognizing that these changes in the AI community are both natural and inevitable. Just as the emergence of commercial computer technology ultimately enhanced the quality and quantity of computer science education and research, this new interest in applied artificial intelligence will in the long run attract more students, faculty, and funding into basic AI research.

In the short run, however, there will be a brief period of instability, when commercial and academic segments appear to compete for resources and dominance. Once the commercial segment has matured, it will have the resources and motivation to support and enhance basic research in the field, just as in any other science associated with commercial uses.

I believe it is in the best long-term interests of the AI community to facilitate and plan for this transformation. Fundamental to this process is to accept two principles: that the interests of the new constituency are legitimate; and that the commercial segment will ultimately come to be viewed by the public as dominant.

What we are experiencing is a natural growth of artificial intelligence from its academic roots to productive application within society. On balance, the effects of this growth can only prove positive for the AI community and society at large. As scientists, we must be prepared to face the new responsibilities and challenges that this transition brings.