Artificial Intelligence: A Perspective[1]

Patrick H. Winston

The primary goal of Artificial Intelligence is to make machines smarter. The secondary goals of Artificial Intelligence are to understand what intelligence is (the Nobel laureate purpose) and to make machines more useful (the entrepreneurial purpose). Defining intelligence usually takes a semester-long struggle, and even after that I am not sure we ever get a definition really nailed down. But operationally speaking, we want to make machines smart.

The typical big-league, artificial-intelligence laboratory, and there are many of them now, will be involved in work like that shown in Figure 1. We at the MIT Artificial Intelligence Laboratory work in robotics, a field spanning manipulation, reasoning, and sensing. We do research in learning, language, and what some people call expert systems, something that I prefer to call design-and-analysis systems, by virtue of the common misuse of the term *expert systems.* We are also involved in issues basic to Computer Science, such as programming and computer architecture.

The Past: Six Ages

The history of Artificial Intelligence can be divided into a variety of ages, as shown in Figure 2. First is the prehistoric time, starting in 1842 when

[1]This paper is the introduction to *The AI Business: Commercial Uses of Artificial Intelligence*, Patrick H. Winston and Karen A. Prendergast, editors, MIT Press, Cambridge, MA, 1984. It is reprinted, with minor revisions, by permission of the publisher.

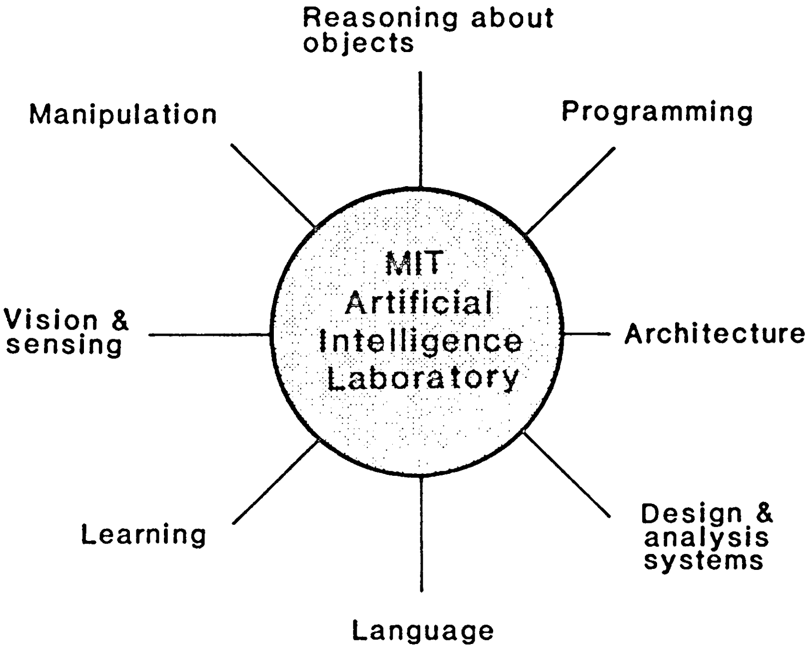

Figure 1. Subfields of Artificial Intelligence.

Charles Babbage first tinkered with his machines. Lady Lovelace, for whom the ADA programming language is named, was Babbage's main sponsor. She was besieged by the press, wondering if Babbage's machines would ever be as smart as people. At that time, she intelligently denied it would ever be possible. After all, if you have to wait for a hundred years or so for it to happen, it is best not to get involved.

The prehistoric times extended to about 1960 because the people who wanted to work on the computational approach to understanding intelligence had no computers. Still, people like Claude Shannon and John von Neumann made many speculations.

Around 1960 we start to speak of the Dawn Age, a period in which some said, "In ten years, they will be as smart as we are." That turned out to be a hopelessly romantic prediction. It was romantic for interesting reasons, however. If we look carefully at the early predictions about Artificial Intelligence, we discover that the people making the predictions were not lunatics, but conscientious scientists talking about real possibilities.

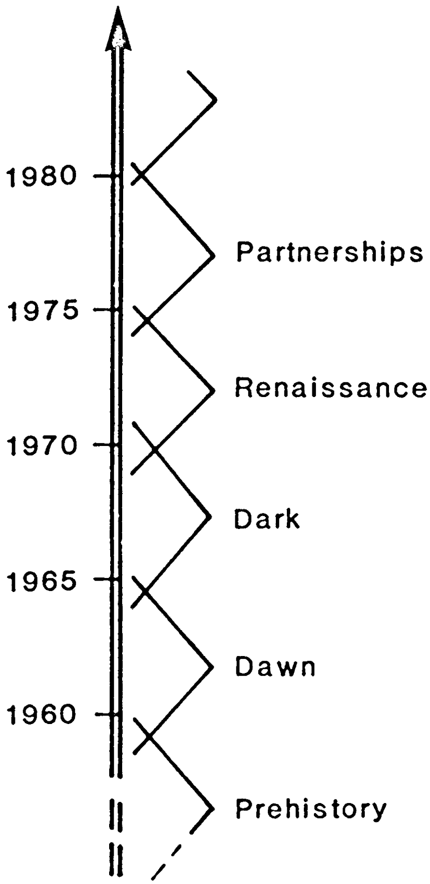

Figure 2. Ages of Artificial Intelligence.

They were simply trying to fulfill their public duty to prepare people for something that seemed quite plausible at the time.

The Dawn Age was sparked by certain successes. A program for solving geometric analogy problems like those that appear on intelligence tests was developed. Another was a program that did symbolic integration, spawning today's MACSYMA and other mathematics manipulation systems. These two examples, integration and analogy, are particularly worth noting because they introduced ideas that have become extraordinarily popular in the creation of expert systems. Retrospectively, the analogy program was based on the paradigm of describe-and-match, and the integration program was based on the paradigm of if-then rules.

I call the next period the Dark Period because little happened. There was a dry spell because the tremendous enthusiasm generated by the Dawn Age made everyone think that the enterprise of creating intelligent computers would be too simple. Everyone searched for a kind of philosopher's stone, a mechanism that when placed in a computer would require only data to become truly intelligent. The Dark Age was largely fueled by over-expectation.

Then we had a Renaissance. During this Renaissance people doing Artificial Intelligence began to make systems that caught people's eyes. MYCIN and other systems developed during this period are the harbingers of today's excitement.

The Renaissance was followed by the Age of Partnerships, a period when researchers in Artificial Intelligence began to admit that there were other researchers, particularly linguists and psychologists, with whom people working in Artificial Intelligence can form important liaisons.

I like to call our present age the Age of the Entrepreneur.

If there were substantial ideas about how to do impressive things as early as 1960, why has it taken until the 1980s to talk about how Artificial Intelligence might be commercialized?

The Successes

Let us agree that something has to be well known and in daily use to be successful. By this definition, there are only a handful of successful systems clearly containing artificial-intelligence technology.

One of the most conspicuous successes is the XCON system (also known as R1) developed by Digital Equipment Corporation and Carnegie-Mellon University for doing computer configuration. Others are DENDRAL and PUFF, products of Stanford University, developed for analyzing mass spectrograms and for dealing with certain lung problems. Still others include General Motors' CONSIGHT system and Automatix's AUTOVISION^{R} II, both of which endow increasingly intelligent robots with a limited but important ability to see.

Other successes are less domain specific. One, a product of Artificial Intelligence Corporation, is INTELLECT, a natural language interface system. Another is MACSYMA, a giant system for symbolic mathematics developed at the Massachusetts Institute of Technology and marketed by Symbolics, Incorporated.

As I recently went over this list of successes with some friends, one pointed out that I had left out some of the most dramatic developments of Artificial Intelligence. One is the LISP programming language, a serious by-product of Artificial Intelligence. It is not surprising that the first major spinoffs of the MIT Artificial Intelligence Laboratory were two LISP Machine companies, Symbolics, Incorporated, and LISP Machine, Incorporated. If we go even further back, there are those who would argue that time-sharing was a major development that came out of Artificial Intelligence. Time-sharing is not Artificial Intelligence, but Artificial Intelligence demanded it.

Expert Systems

Human experts specialize in relatively narrow problem-solving tasks. Typically, but not always, human experts have characteristics such as the following: Human experts solve simple problems easily. They explain what they do. They judge the reliability of their own conclusions. They know when they are stumped. They communicate smoothly with other experts. They learn from experience. They change their points of view to suit a problem. They transfer knowledge from one domain to another. They reason on many levels, using tools such as rules of thumb, mathematical models, and detailed simulations.

An expert system is a computer program that behaves like a human expert in some useful ways. Today's state of the art is such that expert systems solve simple problems easily, occasionally explain their work, and occasionally say something about reliability.

Some expert systems do synthesis. XCON configures computers, for example. Other rule-based expert systems do analysis. MYCIN diagnoses infectious diseases, and the DIPMETER ADVISOR interprets oil well logs.

Currently, there are a dozen or two serious expert systems whose authors have commercial aspirations. By dropping the qualifier *serious*, the number grows to a few thousand. The reason is that creating a simple, illustrative expert system is now a classroom exercise in advanced artificial-intelligence subjects. Soon expert systems will be created in elementary courses in computing at the early undergraduate level.

All of this activity has attracted top-management interest, aroused the entrepreneurial spirit, and stimulated investor curiosity. Are the interest, the spirit, and the curiosity misadvised? It is too soon to be sure because few projects have had time to succeed and none has had time to fail.

Nevertheless there are some questions that can be answered, or at least debated. The list includes the following:

- Can today's technology revolutionize whole industries, or can it just deal with isolated, albeit important, targets of opportunity?
- Where are the most susceptible problems: engineering design, equipment maintenance, medicine, oil, finance?
- What are the obstacles to introducing expert systems: finding the right people, working with the existing human experts, getting snared by technically exciting but off-the-mark ideas?
- How hard will it be to build systems that exhibit more of the talents of real human experts?

Work and Play

A work station is a computer system that can be an exciting, productive partner in work or play. To be a good work station, a computer system must offer many features. First, we must be able to talk to the computer system in our own language. For some systems that language must be English or another natural language; for other systems the language must be that of transistors and gates, or procedures and algorithms, or notes and scales. Second, we must be able to work with the computer system the way we want to, not necessarily the way dogma dictates. In engineering design, for example, some people work bottom up; others prefer to work top down; still others work middle out or back and forth. All should be accommodated. Third, the computer system must constitute a total environment. Everything we need should be smoothly accessible through the system, including all the necessary computational tools, historical records, and system documentation. And fourth, the computer system's hardware must be muscular and the graphics excellent.

Some existing work-station products, like Daisy Systems Corporation's LOGICIAN and GATE MASTER, are extraordinarily important in the design of extremely complicated integrated circuits, often containing tens of thousands of transistors. Another work-station-oriented product, Artificial

Intelligence Corporation's INTELLECT, is not so domain oriented. INTELLECT is designed to be a powerful interface between decision makers and whatever data bases they need to work with. While INTELLECT began as a natural language interface, it is becoming the hub of a multitool, multifile information system, with much of the power residing in the parts having no direct concern with English input and output.

Companies such as Daisy Systems Corporation and Artificial Intelligence Corporation may be merely among the first flow of a potential cornucopia. People are developing work stations for such diverse activities as tax planning, chemical synthesis, robot assembly, musical composition, expository writing, and entertainment.

Where are the likely early successes? Key questions in determining this include the following:

- How important is natural language interaction? What does it take to get natural language interaction?
- What constitutes a minimally muscular computer and minimally excellent graphics?
- How important is it for work-station modules to be able to explain what they do? How important is it for users to be able to intervene whenever they want?
- Who can design and build work stations with human-like intelligence? A dozen people? Any computer engineer willing to learn?

Robotics

An intelligent robot is a system that flexibly connects perception to action. Humans are examples of intelligent robots for the following reasons. First, we can see and feel forces. Consequently we can cope with uncertain positions and changing environments. Second, we have graceful arms capable of fast, accurate motion, together with fantastic hands capable of grasping all sorts of objects. Third, we think about what we do. We note and avoid unexpected obstacles. We select tools, design jigs, and place sensors. We plan how to fit things together, succeeding even when the geometries are awkward and the fits tight. We recover from errors and accidents.

In contrast, most of today's industrial robots are clumsy and stupid. For the most part they cannot see, feel, move gracefully, or grasp flexibly, and they cannot think at all. Most of today's industrial robots move repetitively through boring sequences, gripping, welding, or spraying paint at predetermined times, almost completely uninformed by what is going on in the factory. Of course practical robots need not necessarily resemble people. After all, they are built of different, often superior materials, and they need not perform such a wide range of tasks. Nevertheless many industrialists believe there are many tasks that defy automation with anything short of sensing, reasoning, dextrous, closed-loop robots, with human-like abilities if not human-like appearance.

Consequently an increasing number of major corporations are making bold moves. For a while the general pace was slow in the robot-using industries, and outside of Japan there was little rush to accept and exploit the technology produced by Artificial Intelligence. Now the picture is changing. Small companies have been growing rapidly by supplying industry with turnkey products in which vision a productivity-multiplying component. Large companies like IBM are established suppliers with intensive development efforts underway. Where will this new wave of automation go? How far? How fast? While there is little agreement yet, the following questions help focus the debate:

- Why is it relatively easy to build humanless parts-fabrication factories and relatively hard to build humanless device-assembly factories?
- What are the industrial tasks that require human-like sensing, reasoning, and dexterity? Is it better to eliminate those tasks by redesigning factories from scratch?
- What can be done by exploiting special lighting arrangements? How far have we gone with the simple vision systems that count each visual point as totally black or totally white, with no shades of gray?
- Is the robot itself important? Can we improve productivity with robots alone, or must we think instead about improving whole manufacturing systems?

Today and Tomorrow

Finally there is the question of money. Are the venture capitalists ready for Artificial Intelligence? If so, how long will their readiness last? Is current interest just a passing fad?

Will the commercialization of Artificial Intelligence be driven by need-pull or technology-push? Is Artificial Intelligence becoming commercialized because there are problems that desperately need new solutions, or is it

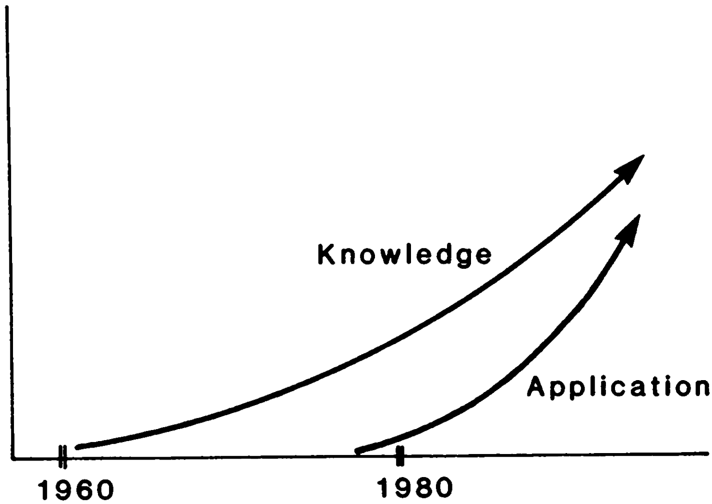

Figure 3. Stiction model of the future.

because there is neglected technology lying around waiting for eager entrepreneurs to make use of it? What sort of progress will there be?

One theory of progress is a kind of mechanical-engineering model, a stiction model, as shown in Figure 3. During the stiction period, the gap between the work in the university research laboratories and the first signs of life in the marketplace constantly grows. Once you get through this stiction period, you move into the period of friction, where the time delay grows smaller, and commercialization marches together with basic research at a steady rate. There are other models of progress. The balloon theory, shown in Figure 4, is one I sometimes believe in when I read the advertising of some of the artificial-intelligence companies. I have a fear that this field has been hyped beyond all belief, and there is a serious danger that it might be oversold.

Figure 5 shows the staircase model of progress. In this model the relationship between the amount of accumulated knowledge and the application of that knowledge is not a linear phenomenon. Knowledge has to accumulate for a long time before there is a sudden burst of entrepreneurial activity that exploits all of the accumulated knowledge. This model says that accumulated knowledge can go only so far and that more knowledge has to accumulate over a period of years before there is another leap for-

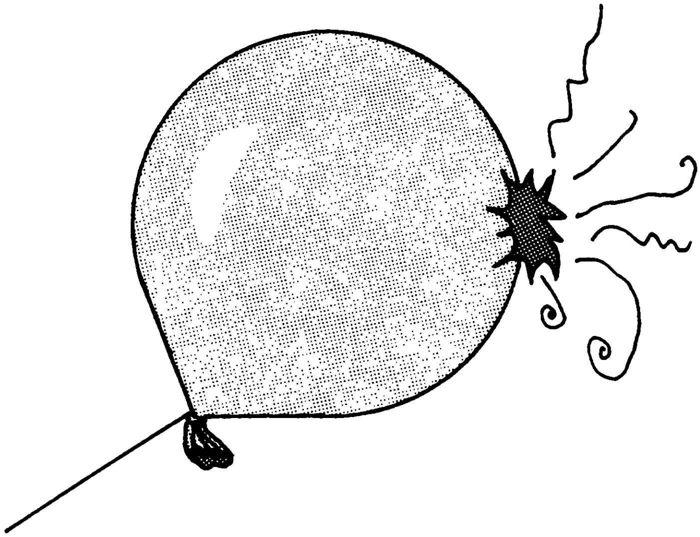

Figure 4. Balloon model.

ward on the applications curve.

We must ask, Which is the correct model for how Artificial Intelligence will develop? At this point, are we merely skimming off the easy problems? Are we repeating in the commercial world what happened in the early days of Artificial Intelligence? If this is a correct model, then we must worry about a Dark Age for the applications of Artificial Intelligence, just as we had one in the basic research area. I do not think there will be a new Dark Age. Too much is happening, as the contributors to this book demonstrate. I believe that the correct attitude about Artificial Intelligence is one of restrained exuberance. It is clear, however, that there are hard-core dissenters on both sides of my position.

For More Information

Feigenbaum, Edward A., and McCorduck, Pamela, 1983, *The Fifth Generation*, Addison-Wesley.

Winston, Patrick Henry, 1984, *Artificial Intelligence, Second Edition*, Addison-Wesley.

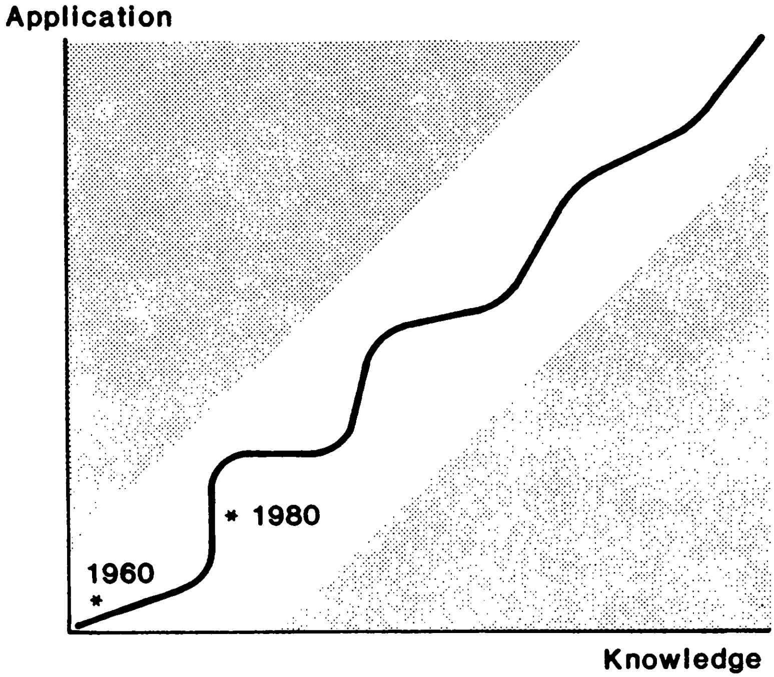

Figure 5. Staircase model.