

# Artificial Intelligence in Transition

Peter E. Hart

*Syntelligence*  
800 Oak Grove Avenue  
Menlo Park, CA 94025

## Abstract

In the past fifteen years artificial intelligence has changed from being the preoccupation of a handful of scientists to a thriving enterprise that has captured the imagination of world leaders and ordinary citizens alike. While corporate and government officials organize new projects whose potential impact is widespread, to date few people have been more affected by the transition than those already in the field. I review here some aspects of this transition, and pose some issues that it raises for AI researchers, developers, and leaders.

THE FIELD OF ARTIFICIAL INTELLIGENCE is in the midst of a deep and irreversible structural change. The older research institutions that were almost alone on the AI landscape (at least in North America) in the late 60's and early 70's have been joined by a host of newer ones; new products based on the fruits of AI research have begun to appear; and the public at large is beginning to believe that "intelligence" can be put in machines.

These changes raise many issues for the AI leadership, and indeed for all who have an interest in the field. New choices must now be made regarding research, development, educational, and business goals. New problems have arisen in communicating both the promise and the present limitations of AI to a wider audience. New institutional problems and opportunities have arisen.

In this article, I will first describe some of the dimensions of the AI transition, and then use these to suggest a number

of issues that deserve close consideration. No panaceas are offered—each individual and each institution will ultimately resolve these issues in their own way—but perhaps enough controversial views will be put forth to stimulate a healthy discussion.

In understanding a review of these issues, I have been struck by their number and the consequent need for selection. The selection made is an attempt to reflect the diversity of questions and choices facing the field as a whole. The reader will have little trouble identifying topics that are of no direct concern, and can pass rapidly to subsequent sections that may be of greater interest.

## Dimensions of the AI Transition

We can get a nice feel for the nature of the AI transition by considering in turn the projects, the institutions, and the people that have comprised the field over the past fifteen years.

**Projects.** First, then, let us recall some typical projects of three different eras: the late sixties, the mid-seventies, and today.

The AI projects of the late sixties were for the most part pure research; *i.e.*, they were "guaranteed useless" in immediate practical terms, though their intellectual contributions were of great importance. As typical examples of that

era we might select the celebrated General Problem Solver developed at Carnegie-Mellon University by Newell and Simon, the SHAKEY robot developed at Stanford Research Institute by Rosen, Nilsson and others, and the blocks world and shape-from-shading vision research pioneered at MIT by Horn and his colleagues.

By the mid-seventies, a new class of applied research projects took its place along side of the continuing stream of basic research projects. Good representatives of this class would be the influential MYCIN medical diagnosis system developed at Stanford by Shortliffe, and the LADDER system for English-language access to distributed databases developed at SRI by Hendrix and others. In addition, a group of "industrial vision" projects were exploring the applicability of earlier vision research to problems of inspection and part manipulation. The hallmark of these projects was their focus on problems of practical significance, though few in this class were brought to the marketplace.

Today, with both basic and applied research flourishing, a new class of commercial products has begun to appear. Examples of these are software products like Intellect, an English-language front end supplied by Artificial Intelligence Corporation, and KEE, a knowledge engineering tool offered by IntelliGenetics. Examples of hardware products are the various Lisp workstations offered by Symbolics, Xerox, and others. In addition, a large number of visual inspection systems are being sold by a variety of vendors. In a sense, the emergence of products like these completes the conventional progression of basic research, applied research, and product development.

**Institutions.** As before, let us recall the institutional landscape of AI in the late sixties, mid-seventies, and today.

In the late sixties, at least in North America, the most prominent AI institutions were a small number of universities and contract R&D organizations; MIT, Stanford, Carnegie-Mellon, and SRI are prime examples. A notable exception to this rule was Xerox, which alone among large corporations (in the U.S.) established a significant AI group at this time.

By the mid-seventies, the traditionally strong AI universities and R&D organizations has been joined by a few more, but there was still very little activity in the private sector.

Today, of course, private sector interest in AI is growing explosively. Within the past few years, any number of major corporations have started significant individual efforts; Schlumberger and Digital Equipment Corporation are good examples. Multi-company consortia have been started in Japan, the U.S., and Europe, of which the Fifth Generation Project has attracted the most attention. In addition, many new companies have been formed to pursue opportunities in expert systems, natural language, and image understanding. Clearly, the big institutional changes of late have been in the private sector.

**People.** It is a bit harder to document the change in the population of "AI people" than it is for AI projects and institutions, but we can get some perspective by considering the influx of people from other disciplines. One measure of

this is the attendance at AAAI and IJCAI tutorials, since these tutorials are typically attended by technologists and managers with little formal training or experience in AI.

The first tutorials, held at the 1980 AAAI meeting at Stanford, were attended by about 300 people. Starting from this base, attendance doubled in 1981 to 600, tripled in 1982 to 900, and in 1983 went up by a factor of six to some 1800 people. Evidently, someone out there is interested in learning more about artificial intelligence.

**A Parallel: Genetic Engineering.** We are told that those ignorant of the past are doomed to repeat it. With this admonition in mind, it is interesting to review developments in genetic engineering, a field whose recent history offers a striking resemblance to our own. Consider some of the parallels:

1. There is a general perception that a genuine revolution is at hand; extraordinary high payoffs are in the offing that go well beyond the normal evolution of a technology. (This perception apparently feeds on the fact that few understand in any detail what the revolution is about.)
2. At the outset of the perceived revolution virtually all the expertise is housed at universities or quasi-university research organizations. For practical purposes, no significant expertise exists in the private sector.
3. Next comes the race to exploit the revolutionary technology. New companies are formed, and large companies establish internal groups. Well over a hundred molecular genetic companies have been started in the past several years. For comparison, it is said that vision companies were being formed at the rate of one a week during 1982. As we have seen, many large companies are now building AI capabilities of some kind. For comparison, today every major drug house world-wide has a gene-manipulation activity of some sort.
4. Finally, philosophical or societal issues are raised concerning the moral or physical hazards associated with the revolution.

I would estimate that the field of genetic engineering is about three years ahead of AI in its evolution. We would do well to watch its development carefully, learn from its successes and try to avoid its mistakes.<sup>1</sup>

### Some Issues

All of these changes present many challenges for those who guide, or who participate in the development of, the field of AI. Since limitations of space force a selection, it seems to me most worth focusing on issues connected with the most radical changes in the field, *viz.*, the explosive growth in

---

<sup>1</sup>Refer to the upcoming article, "Comparing Artificial Intelligence and Genetic Engineering: Commercialization Lessons" for a detailed overview of these two growth industries — the Editor

the private sector. Accordingly, I will consider problems associated with the development of AI products, with interactions between product development activities and research and educational activities, and with allied problems of communication with the vast majority of people who are not AI cognoscenti.

**Issues for AI Development Activities.** Many of the problems associated with the development of AI products are not unique to our field; similar generic problems exist in the development of virtually all software and hardware. However, the field of AI as a whole has thus far had little experience with these problems, and in any case there is value in recalling from time to time a few things that “everybody knows.”

A first basic fact is that products based on AI technology compete in the marketplace with substitute products—*i.e.*, with products based on different technologies that functionally substitute for the AI product. One example of this is the extent to which user-friendly front ends have substituted for natural-language (whether typed or spoken) front ends. Of course, a user-friendly system based, say, on menu selection may be less flexible than a large natural-language system. Further, one could always claim that future natural language front ends will ultimately be the preferred choice. Nonetheless, it is certainly true today that user-friendly systems, arguably based on a non-AI technology, have functionally substituted for AI products in many applications.

Other examples of substitute products can be found in virtually all areas of AI. In manufacturing automation, many problems of orienting and feeding parts can be done more easily by an old device called a bowl feeder than by a vision-based manipulator. Many kinds of classification problems can be solved by conventional statistical techniques more accurately than by a knowledge-based system. Naturally, all approaches have their limitations, and the extremes are usually obvious. (A large robot manipulator can handle parts upwards of a hundred pounds; a bowl feeder for parts that big would be the size of a large building.) Good judgment is required for dealing with the great majority of problems that fall between obvious extremes.

A second point is that AI products frequently do not stand alone. They are used with, or must be incorporated within, other systems in a larger context. For example, an organization interested in providing natural language front ends to databases might in short order discover that it was in the database business as much as it was in the natural language business. Similarly, a group involved with smart manipulators might discover that many of its problems involved visual or tactile sensing. This point and the preceding one suggest that breadth of technical view will be more important for successful applications than for successful research.

A third point is that we frequently underrate the value of simplicity. Many successful AI applications are based on old technology that is considered elementary at the time of its application. For example, many of today’s commercial vision systems use approaches that were pioneered a decade

ago. Some of the most valuable experience accumulated during that decade led to an understanding of which specific techniques could be applied to which specific tasks. Put differently, we learned what does not work, a lesson as valuable as learning what does work.

This leads to the suggestion that we establish, somewhere, a Journal of Disappointing Results, in which investigators could forthrightly describe the limitations of their previously-published theoretical and experimental approaches. Publication in this refereed journal would count as a regular publication; moreover, a subscription to this journal would qualify the subscriber for a 50% discount in SIGCYNIC

Notice, incidentally, that while I have used commercial applications as examples, these observations about AI development activities apply equally well in principle to defense applications.

**Educational Issues.** The AI transition places new stress and demands on the educational system. A full exposition of these issues would properly be a subject for a separate discussion; however, as before, we can gain some perspective by recalling past experience.

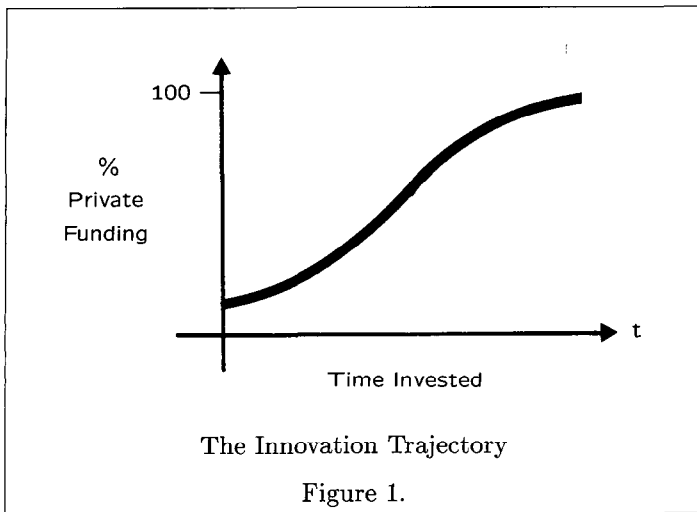
The late sixties were the era of what might be called the Universal AI Researcher. So little had been done that a researcher could attack problems in what today are rather separate sub-disciplines. Moreover, virtually everyone in the field at that time was a migrant from some other discipline; to some extent, AI education at that time consisted of finding such migrants and teaching them LISP.

By the mid-seventies, the educational system was considerably more mature. Interest in applied AI research—together with the emergence of rather clear distinctions between sub-disciplines such as vision and natural language—had an effect on shaping AI education.

Today, as development activities become more prominent, it will be important for the educational system to recognize new needs. Specifically, we can anticipate that AI education will become more like conventional computer science or electrical engineering education, with different curricula and degrees for research, teaching, and development careers.

**R and D Interactions—The Innovation Trajectory.** For at least the past four decades, high technology products have been the end result of what has been called a trajectory of innovation. As the figure suggests, the early part of the trajectory is the basic research phase. Most research is funded by the public sector with relatively low accountability for results, and is performed in university or quasi-university environments. As time progresses and ideas are refined, projects move through an applied research phase toward the product development phase. Private funding becomes far more significant and greater emphasis is placed on achieving workable solutions.

It is interesting to note that this general trajectory has to some extent not been followed by AI in recent years. Some of the most widely cited applied developments have come from universities, while some very significant basic research



results have been obtained in privately funded industrial laboratories. This compression, and even inversion, of the conventional innovation trajectory is one sign of the AI transition. It is not clear that it is a healthy sign, but in any event it is likely to be an anomaly that will disappear in a short time.

It does, however, suggest that more attention be paid to the continuing question of what the appropriate mix is of basic research, applied research, and development projects. The answers will of course be very different for various types of institutions, but one can argue on general grounds that a thoughtful policy has a better chance of being right than a non-policy that exists only by default.

**Communications.** Communications problems have increased in AI as the number of interested constituencies has grown. Again, a backwards look gives some perspective on current problems.

Some fifteen years ago, there were two principal constituencies: those who performed AI research and those who funded it. The funders were public sector agencies such as the Defense Advanced Research Projects Agency, the National Science Foundation, and the Office of Naval Research. These agencies—without which the field of AI would not exist in anything like its present form—often had great technical sophistication, and in fact were staffed in part with people from the field. The relative homogeneity of these constituencies made communication between them easy. (Of course, understanding the other person doesn't necessarily mean you agree with him!)

Today there are many other interested constituencies, and the great variability in their backgrounds and motivations have made communications far more difficult. The performers of AI now include a full range of research, applied research, and development people housed in a variety of organizations. Funders of AI have broadened to include less sophisticated public agencies as well as representatives of the private sector. Moreover, there are now user populations (there was very little to use 15 years ago) who need to be educated and supported. Finally, the public at large has

been made aware of AI and the Fifth Generation, and has (for the moment at least) great interest in the problems and opportunities.

One consequence of the proliferation of constituencies is that communication has become, to some degree, stenographic. Indeed, there are—unfortunately—many instances where content has collapsed to form. For example, a senior executive in a large high technology company may decide to develop “an expert system to do X”, where X might be so unfortunate a choice as “design a major new computer system.” From the executive's point of view there are (human) experts who do X, and since expert systems evidently duplicate human expertise very little further discussion is needed.

As this example illustrates, a particularly delicate communications problem exists at the interface between senior AI technologists and senior operating officials of either public or private organizations. Opportunities for misunderstanding abound, and managing the communication channel is itself a substantial challenge.

A broader challenge centers on how we communicate with the public at large. With the current intense interest in our field, we all need to be aware of our responsibility to communicate honestly and accurately the accomplishments, the potential and the present limitations of AI. The media will be happy to oblige us if we adopt a policy of “better infamous than anonymous.”

## Conclusions

I have tried, by means of illustrations, observations, and parallels to suggest a number of issues that are important for artificial intelligence as a field. Certainly, many other issues could have been brought forth for discussion, but one conclusion is inescapable: artificial intelligence is being transformed very rapidly. Are we ready for the changes?