AI Adventures Worth Writing Home About

Miroslav Benda, Ph.D.
Boeing
P.O.Box 3707, M/S 6A-UJ
Seattle, WA. 98124

Abstract and Introduction

This paper provides a backdrop to my invited talk at the conference. The talk itself will focus on selected successes of artificial intelligence from the industrial perspective. The selected success stories will demonstrate that artificial intelligence has a strong and meaningful influence on our lives by impacting development of products and services we use daily. They will show how the technical results in the field have been used to make a difference in designing complex artifacts, in coordinating our actions, in playing games and in other activities. The emphasis is on transformation processes rather than on specifics of achievements in knowledge representation, case-based reasoning, or sophisticated search techniques. I could not have given this type of a talk ten or even five years ago because this was the initial period of attempts to insert AI into the business processes of companies. Today, ten years later, we have a large number of AI systems that are an integral part of critical business processes, I shall describe some shining examples in my talk. However, the expectations generated ten years ago were much higher than what we were able to achieve in ten years Because of this we have pessimists who talk of failure and "AI winter". Since I want to concentrate on the accomplishments in my talk, I wish to get the negative aspects out of the way here in this paper. Taking the metaphor of a partially filled glass of water, I will describe the empty portion of it in this paper and only remind you of it during the talk. Conversely, this paper just outlines the content of the filled portion while the talk will give you its full taste.

Half empty

> *7 think our mission is to create artificial intelligence. It is the next step in evolution."*

This is how Robert Wright reports (sec [Wright, 1988]) that Edward Fredkin answered his question about the meaning of life. Statements like these make AI vulnerable to criticism since they generate unrealizable expectations especially when the best that AI can produce are systems whose IQ is so small that it cannot even be measured. At the same time these kind of a statements imbue the search for artificial intelligence with a higher purpose and make the activity exciting to the participants and the spectators. The surges of excitement are followed cyclically by backlash. Let me briefly describe the most recent of these cycles, using the Japanese Fifth Generation Project; it clearly illustrates both the half empty and the half full portions of the glass.

The formulation of the Fifth Generation Project was completed by 1981; in April of 1982 the centerpiece of this effort, the Institute for New Generation Computer Technology (ICOT), opened its door. The project had lofty goals (see [2]):

> "... *to cultivate information itself as a new resource comparable to food and energy, and to emphasize the development of information-related knowledge-intensive industries which will make possible the processing and managing of information at will.* "

Some of the statements made to get the project off the ground come close to the Fredkin quote above:

> *"By promoting the study of artificial intelligence and realizing intelligent robots, a better understanding of the mechanisms of life will become possible. ... Mankind will more easily be able to acquire insights and perceptions with the aid of computers.* "

The Japanese were not alone in this enthusiasm. Similar programs that were formulated either concurrently or as a response to the Fifth Generation Project were discussed on national levels the United States and Europe: DARPA's Strategic Computing Initiative, the US Microelectronics and Computing Consortium (MCC), the British Alvey Project and the European Common Market ESPRIT come to mind. The rhetoric in all of these programs contained allusions to the next industrial revolution, global economic interests, national defense strategics and the like. The mass media amplified and disseminated the promises this technology would bring to the world. To give an example, here is a quote from Fortune (see [Alexander, 1982]), a magazine read by the captains of industry:

"The appearance on earth of a nonhuman entity with intelligence approaching or exceeding mankind's would rank with the most significant events in human history."

The examples I just discussed had exposure on the national levels. There were corresponding developments in the leading industrial companies and in the academia. These events helped to start a number of AI companies that struggled mightily to satisfy the demand for a technology with so much promise.

I was very fortunate to be deeply involved in these events: starting in 1983 I was helping to establish artificial intelligence technology at Boeing and I headed the Boeing Knowledge Systems Laboratory till it has successfully completed its mission; more on this later We conducted our business in the context of the statements I related above and these were exciting days The academia and the small companies started by professors marketed their consulting and their products to the government and industry; those of us in the industry tried to adapt the ideas and the products to the needs of the particular industry. These were the days when we thought that solutions to the toy problems of universities would transform our lives

When I look at the situation with ten years of hindsight, we came way short of expectations Looking at proposals that I received from universities for funding of projects, I can only note that capabilities promised to provide Boeing with a competitive advantage by 1990 are yet to be realized I can make similar observations for national level programs such as DARPA's Pilot Associate - where is the electronic copilot on the F-22? Where are the fifth generation machines that "listen, see, understand, and reply to human users"?

For those interested in a comprehensive account of expectations of this period should take a look at E Feigenbaum's and P. McCorduck's book The Fifth Generation (see [Feigenbaum and McCorduck, 1983]) The book made significant contributions to the atmosphere of the period: one of the chapters is called Scuba Gear for the Mind and talks about scientist beginning to create " a sort of scuba gear that will permit the human mind to go where it hasn't been able to go before". They also wrote (p. 125) that "The Fifth Generation project is hard ... but ... TEN YEARS IS A LONG TIME!" (their caps). I was with Ed in Tokyo at the second Fifth Generation conference in 1984 and ten years indeed seemed plenty; now I know that ten years is a short time and that it has not been enough.

However, I have a positive observation to make about unfulfilled expectations. There is a correlation between expectations and the levels at which the expectations are generated. Let me explain. The levels I am referring to go from the level of an individual researcher, to the level of a department, level of a small company, large company, industry, nation, mankind. At each of the levels, L, there is an expectation, E, of what will be possible, say, ten years in the future. Assume that both the levels and the expectations are quantified by numbers. The observation then is that E is proportional

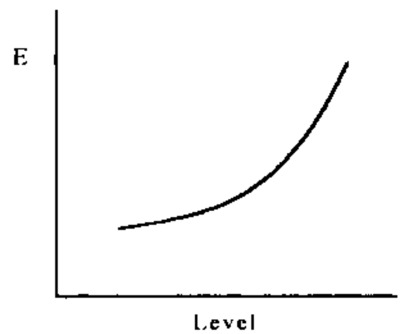

to L^2, $E = cL^2$ (no pun intended)

This implies that the higher the level, the greater the disparity between the expectations and actual accomplishment; in other words, even if the individual researchers did not live up to their promises, their results were considerably closer to reality than the promises made or implied at national levels.

Let me take this as an opportunity to transition to the positive aspects of what transpired in the past ten years. I believe that generating expectations is necessary even if it has negative effects, there probably is some social law that demands the generation of inflated expectations for large scale efforts: the motivating force to cause great changes must be correspondingly large. That force causes alignment in a direction at the higher levels and is propagated to the lower levels where real progress happens; that progress constitutes the filled portion of the glass

Half full

The half full portion of the glass consists of AI accomplishments that have already transformed our lives This will leave out a number of technical achievements with great potential that are close to the point of "condensation". Some of them come from the programs I discussed in the preceding section; to provide a balanced view, let me mention a few examples.

The Fifth Generation Project did achieve some superb technical results. It is my opinion that the following three developments will, in one form or another, have a significant impact in the future: (1) the highly parallel machines capable of executing 100 million logical inferences per second; (2) PIMOS, the operating system for the machines; (3) Kappa-P, a parallel knowledge-base management system and Quixote, an object-oriented database parallel language. These and other software (parallel theorem provers, natural language processing systems, etc.) are available to the computer science community for further development And while there already exist applications developed on the platform described above - intelligent computer aided design, legal reasoning systems - they are experimental and have not yet been incorporated into the real world

Other programs I mentioned in the previous section have similar successes. ESPRIT has turned out to be the most successful of them. It is now in its Third Framework Programme and the European Parliament is examining The Fourth Framework that would run from 1994 through 1998. European computer companies arc satisfied with ESPRIT; there is an estimate that 20% of their information technology products contain technology obtained from their participation in ESPRIT programs. And while DARPA's Strategic Computing Initiative is completed, its results continue to have an impact. For example, some of the concepts and ideas from the Pilot Associate program were used to build an electronic checklist for pilots of commercial aircraft; one such checklist is being tested by airline pilots in an engineering simulation cab at Boeing. The results show significantly fewer pilot errors under complex emergency conditions. It still needs a few years to mature to make a difference but I am confident that no more than five IJCA1 conferences from now I will be able to include increased airplane safety as one of the successes of AI.

Two comments before we sample the filled portion of the glass. First, my definition of an AI system is operational: it is a system that can assist or even substitute for humans in manipulating frameworks of symbols. Let me elaborate. The progress of our civilizations is largely, if not exclusively, due to the use of increasingly powerful systems of symbols, pictures, numbers, diagrams, alphabet, etc.. The advances in technology and in organizational structures can be clearly correlated with our ability to invent and exploit such systems: Mesopotamia and numbers, Egypt and hieroglyphs, Greeks and the alphabet, Arabs and the decimal system, Renaissance and printing, etc. Symbols and the associated mechanism of abstraction enable us to cope with complexity. Tools that help us to shape the world of symbols and the knowledge they carry are what artificial intelligence is about.

The second comment deals with transformations that change established ways. Transformations of this kind typically arise as a result of combining developments in diverse areas: a car or a television is a result of developments in a number of technologies. So, in listing AI successes we must be broad enough to include transformations in which artificial intelligence was one of many technologies responsible for the transformation. I divide the successes into the following classes:

Adopted successes: *accomplishments from AI "prehistory" that turned out to be essential building blocks for AI*

Conceptual successes: *concepts, ideas and goals of AI that were developed into general techniques for building and interacting with information systems*

Theoretical successes: *real world implementations of elements from AI foundations (the formulation of general problem solving, the satisficing principle, heuristic search methods)*

Information appliance successes: *AI systems embedded in run-of-the-mill software packages*

Processes changing successes: *AI systems integrated within large applications supporting major business processes*

My talk will concentrate on the last category. Examples of successes in the other categories are:

Adopted successes: I like to think that it all started with a Cro-Magnon woman who, one day a long time ago, tied a knot on her handkerchief and thus invented a one-bit memory device. Not exactly a personal digital assistant but definitely a step in that direction: a clear example of using symbols to conquer the biological constraints of storing and processing information. Other, much more recent examples, include Aristotelian logic. Boolean algebra, and the predicate calculus. These concepts were later realized in computer languages such as LISP and Prolog. Most of these developments permeate our lives to such an extent that they are taken for granted.

Conceptual successes: The attempts in the fifties and the sixties at making computers to play games, solve symbolic puzzles, and understand language demonstrated that the mechanism of abstraction could be effectively imitated by computers. This had a great impact on the development of programming languages; two of these languages, INTERLISP and SIMULA, originated the object-oriented paradigm, the method of choice for modern development of software. With these languages came the need to operate with symbolic structures much more complex than numbers or letters; this necessitated the development of user interfaces capable of displaying and interacting with information closer to our way of thinking: more parallel, at a higher level, and more intuitive. This type of an interface, windows, pull-down menus, mouse, has become almost universal now.

Theoretical successes: The fact that computer programs can now play chess better than but a few dozen humans also has its origins in AI. Chess and other games spurred the formulation of operators on problem spaces and the development of heuristic search algorithms. These and subsequent refinements together with the advances in building specialized chips constitute the foundation of today's championship machines. The machines, relying on a hardware implementation of search, do not "think" in terms envisioned by AI pioneers, but they do rely on the basic concepts of heuristic search. They have changed lives of chess players: the millions of people who follow chess columns now routinely read game analyses "by" Deep Thought or are replaying its games.

Information appliance successes: I am writing this at the end of March; this is the period when people in the United States are preparing their taxes It is now my fourth year that I am using, together with hundreds of thousand of others, an expert system to do this chore; the system is really an expert in this field and it completely changed the way I prepare my tax return I ma very satisfied with this expert systems. Another AI system that I am using is a grammar checker. Grammar checkers are results of AI exploration into understanding natural language; they are now routinely bundled with word processing packages. There are a number of other systems running on PCs built using expert system methodology. My experience with the grammar checker is less positive than my experience with the tax advisor, probably because I know more about language than about taxes.. On the whole though, it is satisfying to see the productive convergence of personal computers and expert systems and other AI methodologies It shows that even a little bit of intelligence can go a long way.

With this as a background, I am now ready to outline what I will focus on in my talk Our lives and work are enormously influenced by a relatively few corporations; the companies whose products and services we use in our everyday activities and, of course, the companies where we work. 1 want to show the difference AI made and is making for some of these businesses. The question was which ones to choose. I came up with several selection criteria; one of them was to simply take the 30 companies of the Dow Jones Industrial group. After some reflection I decided to select from companies whose products or services would, one way or another, impact my trip from Seattle, my hometown, to Chambery, the site of this conference. Out of the many companies fitting this criterion I chose American Airlines, American Express, AT&T, Boeing, Intel, Microsoft, and Xerox. Each of these companies is effectively using systems integrating artificial intelligence techniques in the execution of its business. I shall describe one example of such system from each of the companies: its purpose, its evolution, and its impact on the business processes The progress is here and it is real.

Conclusion: I showed that the AI glass is definitely not empty as some would have us believe. The successes are real and those who contributed to them have every reason to celebrate In the introduction to this paper I stated that there were essentially no achievements for giving a talk of this type ten years ago. When I look to the future, I would expect that, because of the momentum generated, I would be able to report ten times as much progress ten years from now than was made in the past ten years And if it takes longer, so be it; sooner or later reality will catch up with our expectations!

References:

[1] R. Wright, Three Scientists and their gods: looking for a meaning in an age of information Times Books, 1988.

[2] Proceedings of the International Conference on Fifth Generation Computer Systems, New York, Elsevier-North Holland, 1982.

13] T. Alexander, Teaching computers the art of reason. Fortune, May 17, 1982

[4] E. A. Feigenbaum and P. McCorduck, The Fifth Generation, Addison-Wesley, 1983.